Risk-Informed Incident Management for Nuclear Power Plants

by

Curtis Lee Smith

B.S., General Engineering (1988), Idaho State University
M.S., Nuclear Engineering (1990), Idaho State University

Submitted to the Department of Nuclear Engineering
in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2002

Signature of Author

Department of Nuclear Engineering
August, 2002

Certified by

George E. Apostolakis
Professor of Nuclear Engineering
Thesis Supervisor

Certified by

Michael W. Golay
Professor of Nuclear Engineering
Thesis Reader

Accepted by

Jeffrey Coderre
Chairman
Department Committee on Graduate Students

Risk Informed Incident Management for Nuclear Power Plants

by

Curtis Lee Smith

Submitted to the Department of Nuclear Engineering on September 4th, 2002,
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

Abstract

Decision making as a part of nuclear power plant operations is a critical, but common, task. Plant management is forced to make decisions that may have safety and economic consequences. Formal decision theory offers the potential for a structured approach capable of taking into account risk-related aspects (plant and worker safety, for instance) and, at the same time, important factors like economics and regulatory requirements. Since power generation involves large capital and operational costs, making the decision process more efficient can lead to significant economical savings. With millions of dollars at stake, it is imperative that operational decisions be made in a logical and consistent fashion. In addition to the monetary concerns, a primary driver for this work is the desire to make defensible decisions. Within a structured organization like a nuclear power plant, a variety of interactions take place between groups of decision makers. These groups are asked to provide guidance on a variety of issues, ranging from complex regulatory requirements to planning maintenance activities of standby equipment. By providing an integrated package for decision making, it is believed that tools like the plant risk assessment can be used in a defensible manner as part of the day-to-day operation of the facility.

The goal of this report is to describe a decision methodology for nuclear power plant incidents. Here, incidents are categorized as plant upsets that are not serious challenges to plant safety, but nonetheless require an appropriate response. As part of this decision methodology, risk assessment, worker safety, economics, preferences, and formal decision making models make up the foundation. We describe the construction, analysis heuristics, and inherent uncertainty of these models. From this methodological framework, we developed a prototypical on-line advisory tool that provides decisional advice relevant to incident management. The capabilities of this prototype are discussed along with a demonstration via case studies.

Thesis Supervisor: George E. Apostolakis
Title: Professor of Nuclear Engineering

Contents

Acronyms

AHP	Analytic Hierarchy Process
CAIRS	Computerized Accident/Injury Reporting System
CCDP	conditional core damage probability
CDF	core damage frequency
CVCS	chemical and volume control system
DET	dynamic event tree
F-V	Fussell-Vesely
LERF	large early release frequency
NRC	Nuclear Regulatory Commission
PRA	probabilistic risk assessment
PSA	probabilistic safety assessment
PWR	pressurized water reactor
RAW	risk achievement worth
RRW	risk reduction worth
SDP	significance determination process
SSC	system, structures, and components

Acknowledgements

To my research advisor, Dr. Apostolakis, for always being one step ahead and offering a guiding light through the choppy waters known as "thesis writing."

To the Idaho National Engineering and Environmental Laboratory, under contract to the U.S. Department of Energy, for providing both financial and technical support of this study.

To my Grandparents, for guiding my work and protecting my family. I would especially like to offer gratitude to my Grandmother Elaine for teaching me the joy of chess and cribbage.

To my Parents, Rita and Dennis, for the love, support, and focus they have provided throughout my life – they truly are the most caring people I know. Special thanks go to my father for showing me that not every poker hand is a winner, but a good poker player can still win.

To my children, Drew, Cory, Seth, and Gabrielle, for reminding me that imagination, creativity, and love are what matters most.

And lastly, to my wife Jacquie, for teaching me the joys of Chemistry and Calculus and for always being there as my partner, forever and ever, Amen.

Curtis

Executive Summary

In this thesis, we have outlined the framework behind formal decision making for nuclear power plants, leading to the development of an Internet-based incident management advisory system (simply called "the prototype"). The goal of this prototype is to facilitate selection of a preferential decision alternative in response to an incident and provide technical justification for the decision. During our research, we encountered a variety of technical issues, including the determination of decision maker preference functions, consistency between decision performance measures, automation of a decision model, and the solution of time-dependent decision processes.

Incident management, the decision making that follows events ranging from non-safety-related component outages to complex transients, should be based upon decision science as characterized by the game theory of von Neumann and Morgenstern (von Neumann and Morgenstern, 1944). While all plants have well-defined operating procedures which defend against serious plant challenges, it is of more concern to us to investigate incidents since they provide a "gray" area of plant operation – an area where latitude is provided the plant operators to consider alternatives in response to the incident. By combining probabilistic models of plant operation with formal decision theory, we have the potential for a structured approach capable of taking into account risk-related aspects (plant and worker safety, for instance) and, at the same time, important factors like economics and regulatory requirements.

An obvious motivation behind the desire to have an incident management system is money. Since nuclear power plants have both a large capital cost and significant operational costs, making the decision process more efficient can lead to potentially large economical savings. But, in addition to monetary concerns, an important driver for our work is the desire to make defensible decisions. Within a structured organization like a nuclear power plant, a variety of interactions take place between groups of decision makers. For example, the people responsible for the facility risk model frequently are asked to provide guidance on a variety of issues at the plant. By providing an integrated methodology for decision making, groups such as the plant risk analysts will be able to facilitate use of their activities, in this case the probabilistic risk assessment (PRA), in a defensible manner. Further, they will be able to define a proper context for PRA application within the framework of decision making, thereby promoting the use of risk technologies as part of the day-to-day operation of the facility.

A natural question then, at this point, is "what exactly are incidents at a nuclear power plant?" The types of incidents that our prototype is expected to address include component failures such as pumps, valves, or instrumentation; degradations such as leaks from coolant systems; and potential system impacts such as non-compliant components. But, as we discuss in the thesis, the approach to decision making has been generalized in the prototype such that a wide variety of events can be modeled. This generalization does not imply that the prototype will handle every situation with aplomb, but it is hoped that the methodology does encompass a large number of potential incidents.

ES-1. The Prototype

In order to bring together the decision modeling and analysis heuristics, we have developed a decision advisor prototype. This prototype is used as a test bed for new ideas and, ultimately, will demonstrate the decision-making technology. During the development of the prototype, we focused on four major aspects:

1. A primary controlling module to collect incident information and subsequently determine the decision model.
2. Preference models representing the decision maker's beliefs (via a value tree and associated disutilities).
3. Supporting analysis modules (for example, the PRA, economic models, and worker safety models).
4. An analysis module to solve the decision process.

These four aspects make up the structure of the prototype. From these, we constructed a "activity tier" diagram (shown in Figure ES-1) which defines the information and analysis flow framework. Within this framework, a total of five stages are represented. Each stage embodies a unique portion of the overall process of determining a preferential decision option from a list of alternatives.

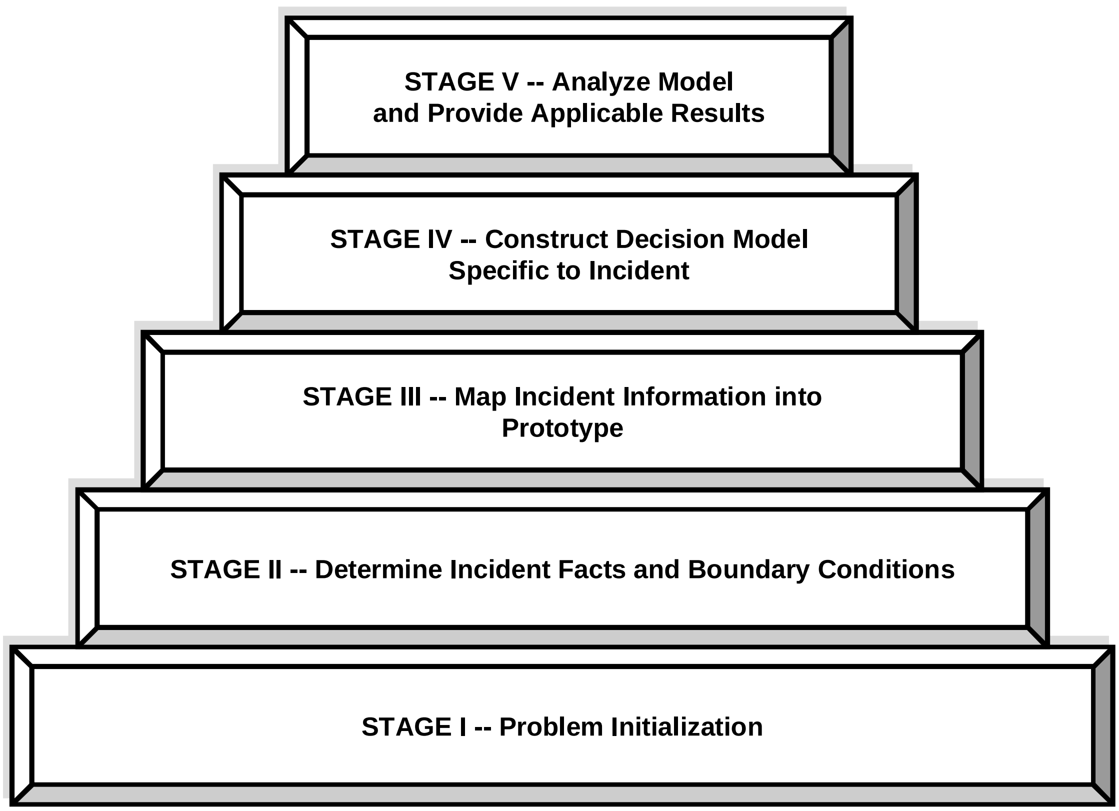

Figure ES-1. Framework tiers embodied in the decision advisor prototype.

Let us briefly describe the stages. First, Stage I of the prototype represents the initialization of the decision process. The current version (version 1) of the prototype uses a web-centric approach where the prototype runs off a web server. During this stage, the user is presented with the value tree, performance measure weights, and associated disutilities. We will discuss the technical basis for these items later in this section. An example of the graphical user interface for the prototype is shown in Figure ES-2. Note that the prototype is multi-lingual, where the language can be changed by selecting from a list of available options.

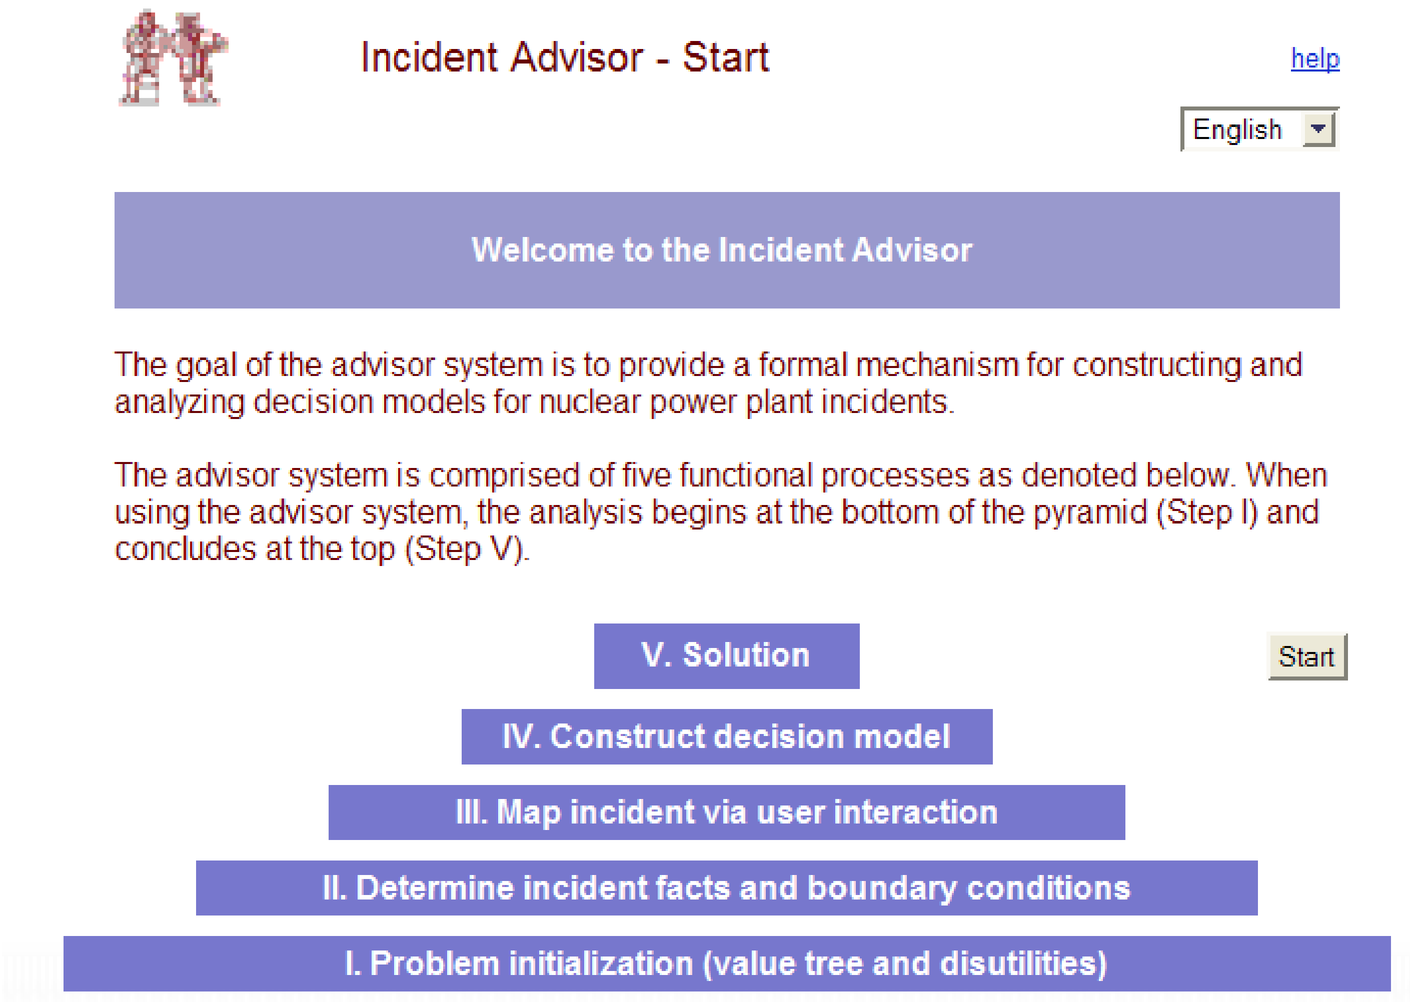

Figure ES-2. Example of the user interface for the decision advisor prototype.

In Stage II of the prototype application, the user must provide incident-specific facts and related boundary conditions. For example, the types of data that must be supplied by the user at this stage include:

- The type of incident (component or initiator related)
- The current reactor state
- The time until the next scheduled outage
- Impacts to plant operations through component degradations

Within this stage, the user will provide a variety of information specific to the incident. Some of this information may cause the prototype to ask the user for additional information (for example, if a component degradation impacts initiating events, the user will be queried for this information) in later stages.

In Stage III, the details of the incident are entered into the analysis data stream. The knowledge base contains a variety of potential decision alternatives, options like continuing as-is, shutting the plant down to fix the problem, repairing the problem at power, etc. Incident conditions that are tied to a decision (through the knowledge base) will cause that decision to automatically be used in the analysis.

Once the relevant decision information is entered into the prototype, the user enters Stage IV. Here, the prototype will automatically construct the decision model that will be analyzed to determine the preferential decision option. Prior to the actual analysis, the prototype displays the information that has been collected from the user in order to offer a final check of the decision model inputs. But, once the user is satisfied with the analysis conditions, the decision model is constructed by the prototype. The process that takes place is to evaluate a generic influence diagram/decision tree via a static, sequence-based "roll-back" calculation (Clemen, 1996). We have also implemented a novel simulation-based analysis module that envelops the entire decision process with a discrete event simulation framework. Both the "roll-back" and simulation modules have been embedded in the prototype without the reliance on third-party analysis tools.

Once the point estimate results are shown on the screen, the user has entered Stage V of the decision advisor prototype. Stage V is the final step in the analysis process, whereby the user is allowed to view the results. The user is shown the decision alternatives ranked by preference along with a numerical score calculated by using an additive form of multi-attribute utility theory:

$$\begin{aligned} PI &= w_{economics}\ u(economics) + w_{dose}\ u(dose) + w_{accidents}\ u(acc.) \\ &\quad + w_{safety}\ u(safety) + w_{stakeholders}\ u(stakeholders) \\ &= PI_{economics} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{stakeholders} \end{aligned} \qquad \text{(ES-1)}$$

where w_i is the weight of the *i'th* performance measure and *u( )* is the disutility associated with a particular outcome of the *i'th* performance measure. While non-linear forms of PI are available, they are generally not utilized. As noted by Clemen (1996) "…in extremely complicated situations with many attributes, the additive model may be a useful rough-cut approximation."

Note that the PI shown in Equation ES-1 is only a point estimate. This value is only used to give the user a "feel" for the decision alternative outcomes. For decision making, one must rely on the expected value of the PI. Consequently, the prototype performs Monte Carlo sampling on the input variables to PI as part of the expectation calculation. But, once the PIs are known for the decisions, we rely on the decision rule to select a preferential decision. In our case, we use disutility functions, which implies that we desire to minimize negative outcomes. Thus, decisions with low PI are preferred, or:

$$Decision(1) = \min(E[PI]_i) \qquad \text{(ES-2)}$$

where $Decision(1)$ is a the preferred decision alternative and $E[PI]_i$ is expected value of the PI for the *i'th* decision alternative.

When running the Monte Carlo uncertainty analysis, we are estimating the expected value of PI for each decision alternative, or:

$$E[PI] = \int_0^{\infty} \left(\sum_{i=1}^{5} w_i \, u(x_i \,|\, y_i) \right) \pi(\bar{\mathbf{x}}_i \,|\, \bar{\mathbf{y}}_i) \; d\bar{\mathbf{x}}_i \qquad \text{(ES-3)}$$

where i indicates the performance measure; w_i is the weight of the measure; u_i is the disutility for the measure; $\pi(\ )$ is the epistemic uncertainty over the decision analysis parameters and models; x is the decision analysis parameters and models (which, in general, is a vector of factors for each decision alternative and performance measure); and y is the boundary conditions of the problem (i.e., the evidence or the facts).

We exercised the prototype via application of two case studies, first a leaking steam generator tube and second an inoperable pressure transducer. For both case studies, we calculated the point estimate results, performed sensitivity calculations, and ran uncertainty calculations.

ES-2. The Decision Model

The primary focus of the decision model was to utilize a structure drawn from influence diagrams. Within this model, we identified six major parts that are germane to incident management:

- Decision alternatives – these include the options specific to the incident.
- Incident specific elements – these include the possibility for repair, the type of failure mechanisms present, and other unique features related to the incident.
- Boundary conditions – these include the plant state and time until the next outage.
- Plant upsets – these include initiators such as transients/leaks that lead to complications.
- Plant response – these include the plant system response to any upset conditions.
- Outcomes – these include the outcomes of interest to the decision maker.

The decision model embodied in the prototype is a generalized influence diagram. Other analysis modules support the main decision model, and include economic impacts, risk analysis, and worker safety. Interacting with these modules is a knowledge base that supports the generation of an incident-specific decision model. But, once the decision model is constructed, the prototype still needs to analyze the model. Consequently, we developed an appropriate analysis methodology to accompany the decision model.

ES-3. The Value Tree

We collected preference information from our decision makers during two workshops. In the initial workshop, we developed the structure of the "value tree" for incident management (shown in Figure ES-3). When constructing the value tree, deliberation played a critical role both in shaping the structure of the tree and in determination of the performance measure weights. For example, we had one case where the decision makers were divided into a group of three and an individual, where the value tree weights were significantly different between the two groups. Rather than taking a average of the decision makers, deliberation resulted in the group of three modifying their weights to be similar to the individual's weights. As part of this deliberation, we ran "sanity" checks on the results so that the decision makers could see the impact from the discussions.

In the workshops, we utilized a technique known as the Analytical Hierarchy Process (AHP) wherein the decision makers are ask to perform pair-wise comparisons between potential outcomes (Saaty, 1980). First, we used AHP to determine the value tree performance measure weights, specifically in the context of incident management. When we use AHP to determine these weights, we are comparing one entity (the worth of cost) against another entity (the worth of worker safety). Also, these performance measures are within an order of magnitude from one another (one is not 100 times more important than the other in the context of incident management). Further, they do not have arbitrary outcomes (for example, for radiological dose, the focus is "dose" as an impact, not on any one particular value for dose). These conditions on the weights are consistent with those suggested by the developer of AHP (Saaty, 1980; Saaty, 1997).

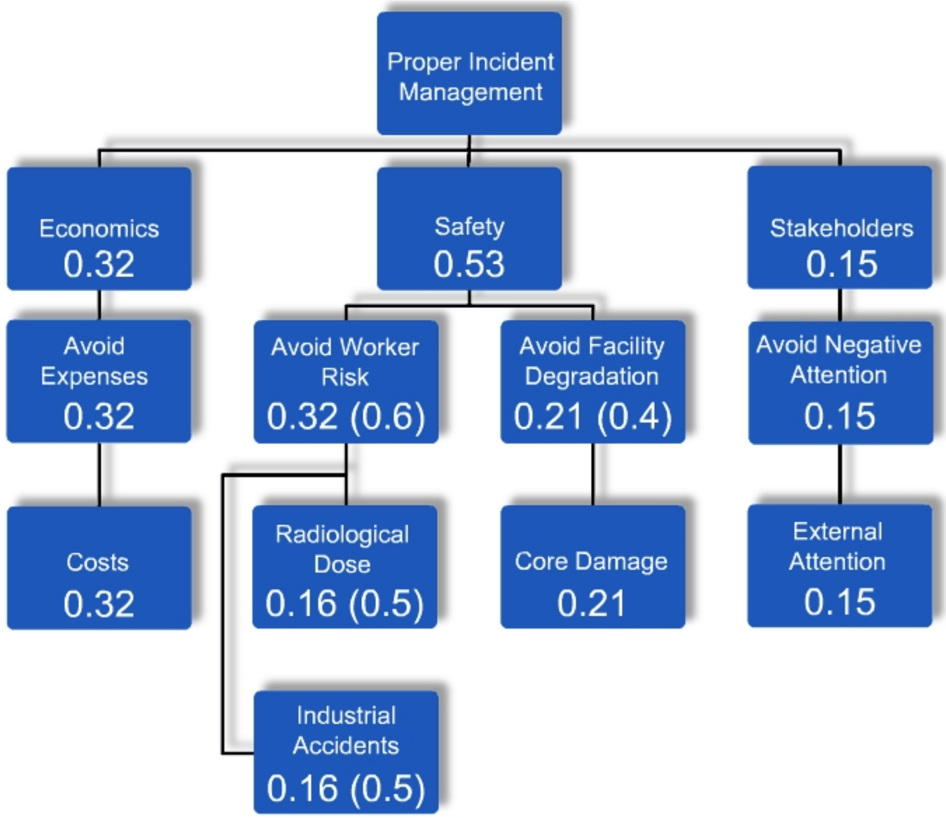

Figure ES-3. The value tree derived from our decision makers for incident management.

Following the value tree weight determination, we used AHP to develop the performance measure disutility functions. Through this process, we were able to illustrate limitations in the AHP-derived disutility functions by utilizing multi-dimensional consistency checks using fixed indifference levels. A surprising result from our analysis was the lack of fidelity when using AHP for disutility determination. AHP is a technique that has seen increasing application in the field of decision analysis. But, we found several technical issues that limited its usefulness.

1. AHP yields disutilities that imply the decision makers are very risk prone.

2. AHP preserves the certainty equivalent only for the low scale regions (e.g., zero to one million euro). In higher regions, the certainty equivalent is suspect. For example, the AHP results suggest that the decision maker would only be willing to pay 20,000 euro to avoid a "50-50" lottery (say a flip of a coin) where the two outcomes are either lose nothing or lose 1 billion euro. In reality, the decision makers would be willing to pay much more than 20,000 euro to avoid large losses of this type.

3. AHP disutility curves are extremely sensitive to small changes in the initial scale regions. In other words, it is possible to obtain dramatically different disutility functions depending on the nuances of setting up the AHP scales.

4. The sanity checks performed on the AHP results showed several inconsistencies. For example, the AHP application indicated that a lethal dose should be approximately 0.2 Sv, when in fact it is approximately 7 Sv.

To overcome the issues related to AHP and disutilities, we developed an approach called "measurable equivalence" in order to construct the required disutility functions. This approach allows us to ensure that the maximum outcome (the worst case) for each performance measure has about the same level of "consequence." Further, we utilize performance measure indifference points in order to constrain the disutility function, where the constraint is made by actual measurable equalities. This second feature is used to bring real data into the decision process while simultaneously helping to reduce the subjectivity present when utilizing preference information. In order to transpose equivalencies from one performance measure to another, we needed to have at least one disutility function fully specified. We chose to determine the cost disutility, where we determined points on this function by way of lottery equivalence questions poised to our decision makers. We then proceeded to determine the remaining disutility functions by applying the measurable equivalence approach.

Examples of the measurable equivalencies that were utilized are shown in Figure ES-4. Note that these equivalencies are treated as indifferent points (similar to indifference curves in decision theory) and are used to obtain individual points on the applicable disutility function. An alternative approach to determining disutility functions would have been to convert everything into monetary values (e.g., euros) by determining a linear conversion factor for each disutility curve. But, we deemed this approach too limited and instead utilized our measurable equivalence approach where individual disutility points are determined and the disutility curve constructed from these points. Note that this approach allows the flexibility to bring in different sources of information, including elicitation and deliberation on the part of the decision maker.

After determining each of the disutility functions, we proceeded to test them to see if they seemed reasonable. Consistency checks were performed using the equivalence-based disutilities – the results of which showed very good agreement with expected results.

To perform the sanity checks between the performance measures, we need to compare two measures against each other at a specified PI value. For example, if we postulate a value of 0.0001 for $PI_{\text{economics}}$ then, in theory, the decision maker should be indifferent to a value of 0.0001 for any of the other measures (since they are additive). In terms of equation ES-1, we are equating individual PI measures:

$$w_{economics}\ u(economics) = w_{dose}\ u(dose) = 0.0001 \qquad \text{(ES-4)}$$

Since we know the weight of economics (0.32) and dose (0.16), we can then determine a cost and dose from the disutility curves that will result in a PI of 0.0001. In this case, we find that the cost is approximately 1.6 million euro and the dose is approximately 6 Sv.

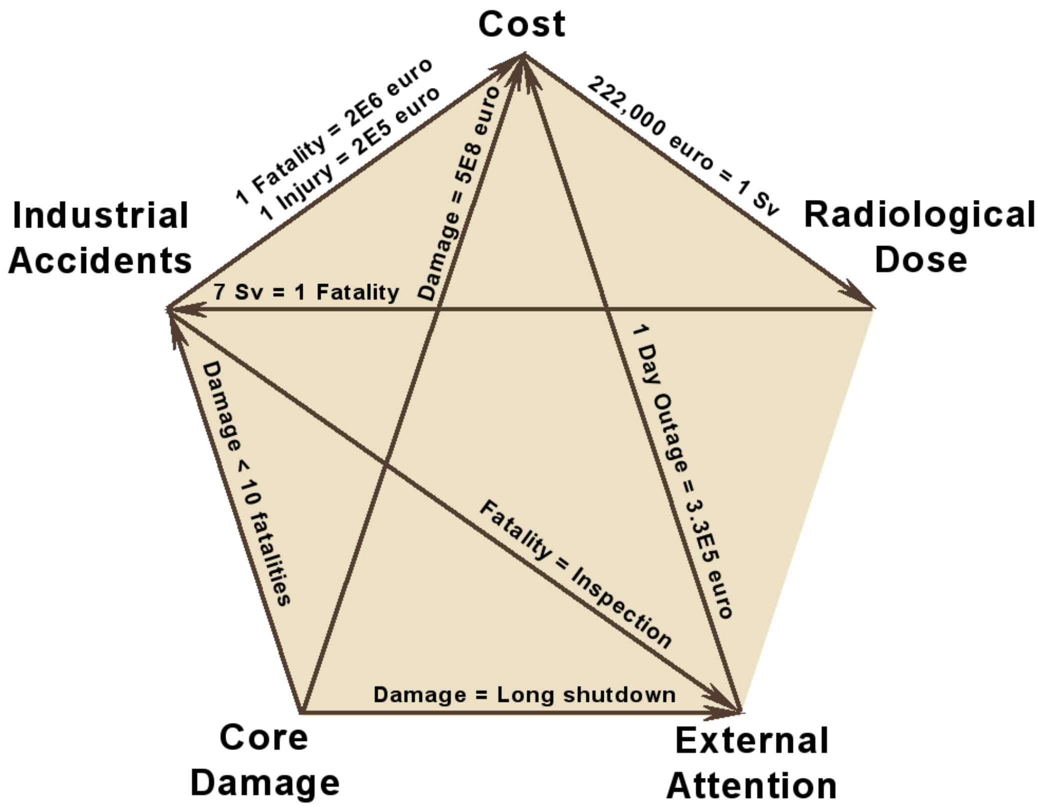

Figure ES-4. Measurable equivalence factors used to develop the final disutility functions. The arrows indicate the "direction" of influence from one measure to another (e.g., receiving a dose of 7 Sv induces a fatality).

We plot the PI comparisons for each of the performance measures in Figure ES-5. Evaluating this sanity check figure, we see that the consistency between performance measures is quite good. For example, we see that a fatality is equal to approximately 2 million euro, which is also equal to a dose of about 7 Sv. Also, a list of equality comparisons is provided in Table ES-1, which contrasts the disutility results obtained by way of AHP and that by measurable equivalence. In every case, the measurable equivalence results are superior to those obtained via AHP.

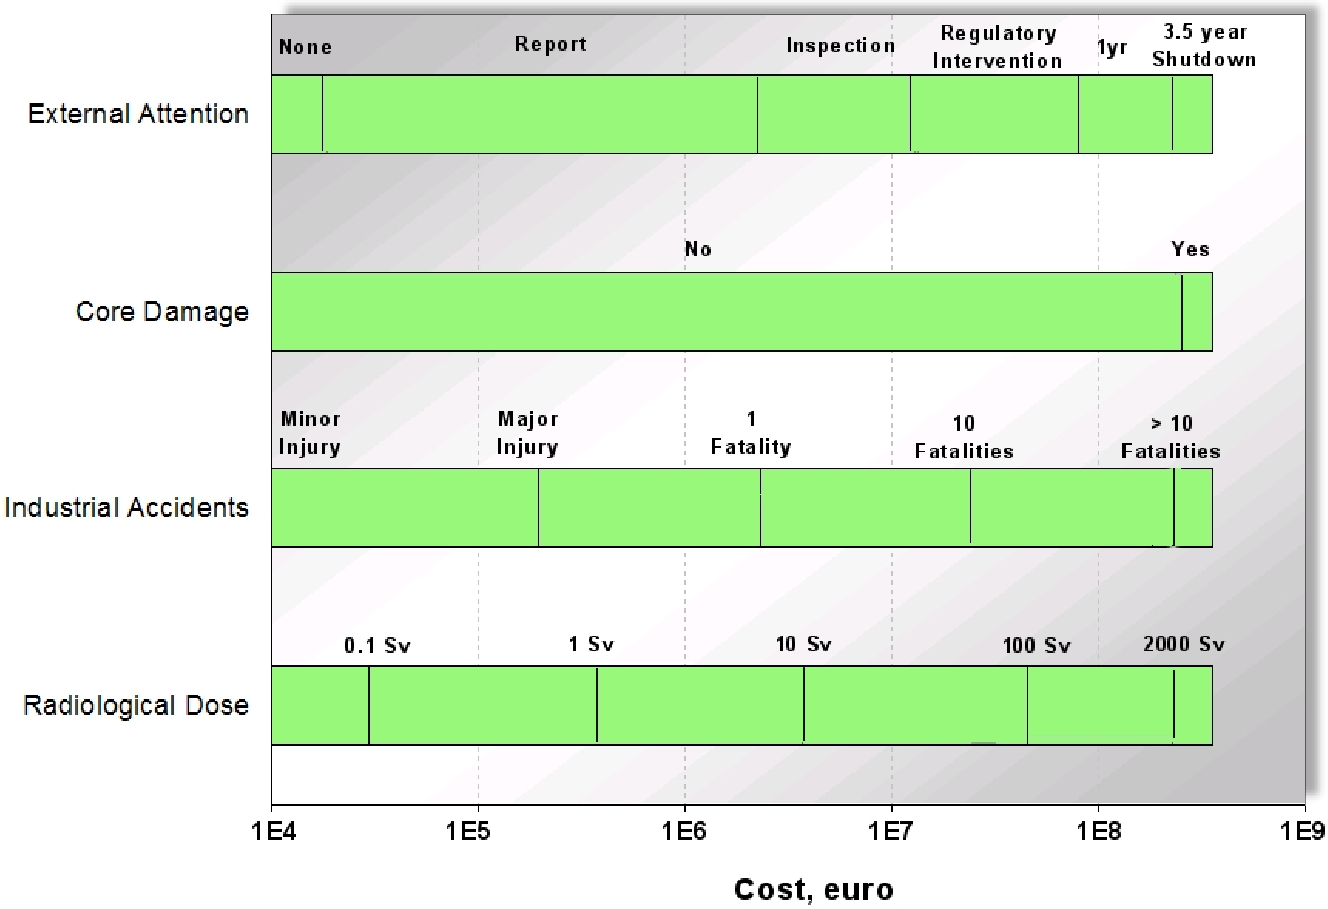

Figure ES-5. Consistency checks plotted as a function of cost, for the measurable equivalence-based performance measures.

Table ES-1. A comparison of consistency checks for AHP-derived and measurable equivalence-derived disutilities.

Consistency Check	Original AHP Results	Measurable Equivalence Results
Fatality cost	10 million euro	2 million euro
Lethal radiological dose	0.2 Sv	7 Sv
Severe injury cost	3 million euro	200,000 euro
Core damage cost	15 million euro	460 million euro
Regulator inspection cost	3 million euro	2.6 million euro

ES-4. Analysis Modules

In order to assist a user confronted with complex nuclear power plant decisions, we designed the prototype to rely on an underlying domain-specific knowledge base. Included in this knowledge base are modules such as:

- Decision maker preferences via the value tree (and corresponding weights) and the measurable equivalence disutilities.
- PRA impacts through initiator upsets and safety system responses.
- Economic, worker safety, and radiological dose deterministic modeling.
- Decision alternatives related to incident management.
- Plant operational state determination.
- Sensitivity and uncertainty treatment.

The major analysis modules are depicted in Figure ES-6, where we also indicate the primary impacts on the value tree that are modeled in the prototype. For example, we utilize a "worker actions" module that depicts the likelihood of a worker experiencing an injury or fatality. If a worker were injured, this outcome would impact the industrial accidents performance measure.

ES-5. Simulating Decision Processes

To assist in the calculation of dynamic events (like what one might find during plant operation in response to an incident), we investigated the possibility of enveloping the entire decision process in a simulation representation. Since decision processes involve stochastic outcomes, we evaluated both traditional static and dynamic models. For treatment of the dynamic models, we utilized a modified version of the Metropolis simulation algorithm (Metropolis et al, 1953), but we extended the routine to encompass

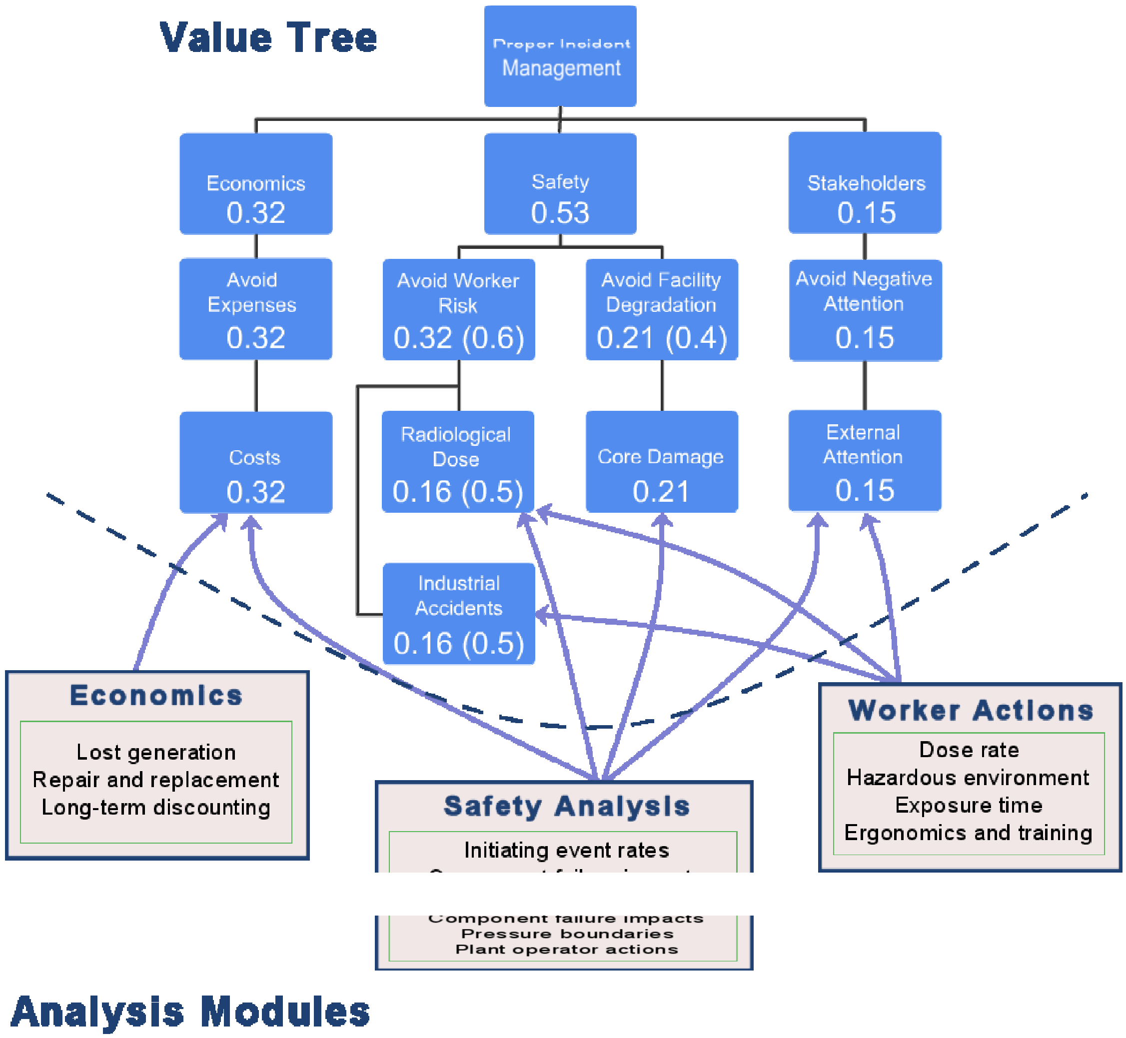

Figure ES-6. The prototype analysis modules and their impacts on the value tree.

decision processes for nuclear power plants. Within our decision framework, we are concerned with plant upsets – via initiating events – and the reliability of safety systems. We found that in some cases, static models will adequately represent the decision model. But, in other cases, static decision trees and fault trees provide only a rough approximation to the exact answer. Note though that static models may be solved very quickly, while simulation-based approaches take much longer. Consequently, we structured the decision advisor prototype such that a two-level analysis approach is possible. A static model can be used to provide decision advice immediately while the prototype continues to process the simulation model.

Within this analysis framework, we described both the deterministic and aleatory models required to assist in the calculation of preferential decision options. In general, the calculation of preference between decision alternatives focuses on the expected PI. Since we are dealing with disutilities for the prototype, we seek to have decision options with low E[PI]. Within the deterministic framework, we determined models specific to plant operation and economic impacts. For plant operations, initiating events were used in two ways, first to determine the probability of getting to a decision-specific state (say a transient) and then second to determine the outcome of that state. For the economics modeling, we utilized the work done by Burke, Aldrich, and Rasmussen on nuclear power plant costs (1984).

The aleatory modeling focused mainly on plant operations and the safety assessment. Within the context of plant operations, we have modeled plant upset conditions as leading to impacts in the PI performance measures such as core damage, cost, and external attention. Further, we determined models for industrial accidents and radiological dose, both of which affect workers.

The goal of the decision-process simulation is to determine potential outcomes, measure their individual impacts, weight the impacts by their respective likelihood, and then calculate an overall E[PI] for each decision alternative. Then, the decision alternative with the lowest E[PI] is considered to be the preferred option. In order to implement the event simulation, we provide the mathematics needed for both a "thinning event" simulation and a "lifetime event" simulation. Thinning event simulation questions the state transition probability at each incremental time step while lifetime event simulation questions the time duration expected in a particular state. In general, the lifetime event simulation approach is much more computationally efficient, especially for the case of reliable components and systems. Consequently, this is the method we implemented for the simulation module. As part of the simulation discussion, we described, step-by-step, the details of the calculation involved in simulating a decision.

ES-6. Conclusions

We have provided a detailed methodology for formal decision making relevant to an incident management advisory system in nuclear power plants. This framework includes the treatment of a variety of deterministic and stochastic models, including decision maker preferences; the plant PRA; economic costs; worker industrial and radiological safety; external attention; and plant state changes. Primary contributions include the encapsulation of the decision methodology in an Internet-based prototype, the discovery of limitations in the AHP for disutility application, the "measurable equivalence" method of disutility determination, and the solution of decision processes via a custom simulation routine.

As the research proceeded, a variety of secondary accomplishments were also realized. Our treatment of epistemic uncertainty was comprehensive and included portions of the decision model such as preference functions and the value tree weights. The discussion of the potential modifications and limitations when using a PRA for decision making points out areas of concern. Since we do use a variety of PRA information as part of the prototype, we developed a XML schema specific to this information in order to facilitate the transfer and manipulation of data structures found in a typical nuclear power plant PRA. We offer this format to the analysis community with the hope that it will encourage the exchange of reliability and safety information.

In closing, this thesis provides the script for decision making, a *play* that casts process models; applicable, informed decision makers; and formal decision-making technologies together with the goal of assisting, not replacing, human judgment. While decisions take place on the stage of uncertainty, it is important to remember that decision science provides a solid foundation for the framework, and associated prototype, described within this document.

Risk Informed Incident Management at Nuclear Power Plants

"Indifference is the essence of inhumanity." — *George Bernard Shaw*

1 Introduction

The management of personnel, operator actions, and decision making during nuclear power plant operations is a critical task and raises specific issues that deserve attention. Decisions are to be made by the plant management that may have both safety related and economic consequences. Attributes like worker exposure, radiation release, bad publicity, and regulatory intervention may conflict with concerns like loss of income due to plant shutdown or reduced power, repair, and maintenance costs.

It has been suggested that probabilistic risk assessment (PRA) insights can contribute to making better decisions in the nuclear industry (Apostolakis, 2000). Furthermore, research has shown that decision makers can frame judgement in a formal decision analysis framework. The combination of PRA and decision theory offers the potential for a structured approach capable of taking into account risk-related aspects (plant and worker safety, for instance) and, at the same time, important factors like economics and regulatory requirements. Incident management, the decision making the follows events ranging from non-safety-related component outages to complex transients, should then utilize this combination of PRA and decision science. While many plants have well-defined standard operating procedures which help to defend against serious plant challenges, it is of more concern to us to investigate incidents since they provide a "gray" area of plant operation – an area where latitude is provided the plant operators to consider alternatives in response to the incident.

Since nuclear power plants have both a large capital cost and significant operational costs, making the decision process more efficient can lead to potentially large economical savings. Examples of the magnitude of costs under consideration include the operations and maintenance budget of the Forsmark site in Sweden around $100 million per year (Kovan, 2002) and the total nuclear power production cost for the U.S. industry of $13 billion per year (Nuclear News, 2001). With millions of dollars at stake, it is imperative that operational decisions be made in a logical and consistent fashion. Consequently, early in the research for this thesis, it was a desire to translate the science of decision making into a prototype advisory system, whereby this system can be utilized in an operational setting. The details of this prototype will be discussed later in Section 5. The types of incidents that this prototype is expected by address include component failures such as pumps, valves, or instrumentation; degradations such as leaks from coolant systems; potential system impacts such as non-compliant components; and tradeoff questions such as inspecting degraded equipment in one plant state versus another.

In addition to the monetary concerns, a driver for the work embodied in this thesis is the desire to make defensible decisions. Within a structured organization like a nuclear power plant, a variety of interactions take place between groups of decision makers. For example, the people responsible for the facility risk model frequently are asked to provide guidance on a variety of issues at the plant, ranging from complex regulatory issues to planning maintenance activities of standby equipment. But, providing risk input to decisions while operating in isolation may lead to inappropriate guidance since interactions between other aspects important to the decision may be missed. By providing an integrated package for decision making, groups such as the plant risk analysts will be able to facilitate use of their activities, in this case the PRA, in a defensible manner. Further, they will be able to define a proper context for PRA application within the framework of decision making, thereby promoting the use of risk technologies as part of the day-to-day operation of the facility.

The goal of this thesis is to describe a decision methodology for nuclear power plant operations and planning, focusing on the management of incidents. As part of this methodology, PRA and formal decision making techniques make up the foundation. This foundation will be capable of taking into account decision maker preferences. In addition, we will utilize the methodological framework to construct a prototypical on-line advisory tool for decisional advice relevant to nuclear power plant incident management. The capabilities and limitation of this prototype will be discussed along with a demonstration via case studies. Note that the scope of the analysis described in this thesis is specific to incidents, which by definition, occur prior to a core melt event. Situations where decisions are made after a core melt event fall into the category of accident management (Catton and Kastenberg, 1998; Jae and Apostolakis, 1992). This distinction of incident-versus-accident management is shown in Figure 1, where minor incidents progress to major accidents from upper-left to lower-right. Since we focus on incident management, we will not address items such as core melt mitigation or emergency response. We will instead evaluate a wide variety of impacts related to plant upsets, ranging from items such as the hourly costs to repair inoperable components; the likelihood of a worker injury or fatality; the occurrence of plant operational state transitions; and the rise of external attention following unwanted events.

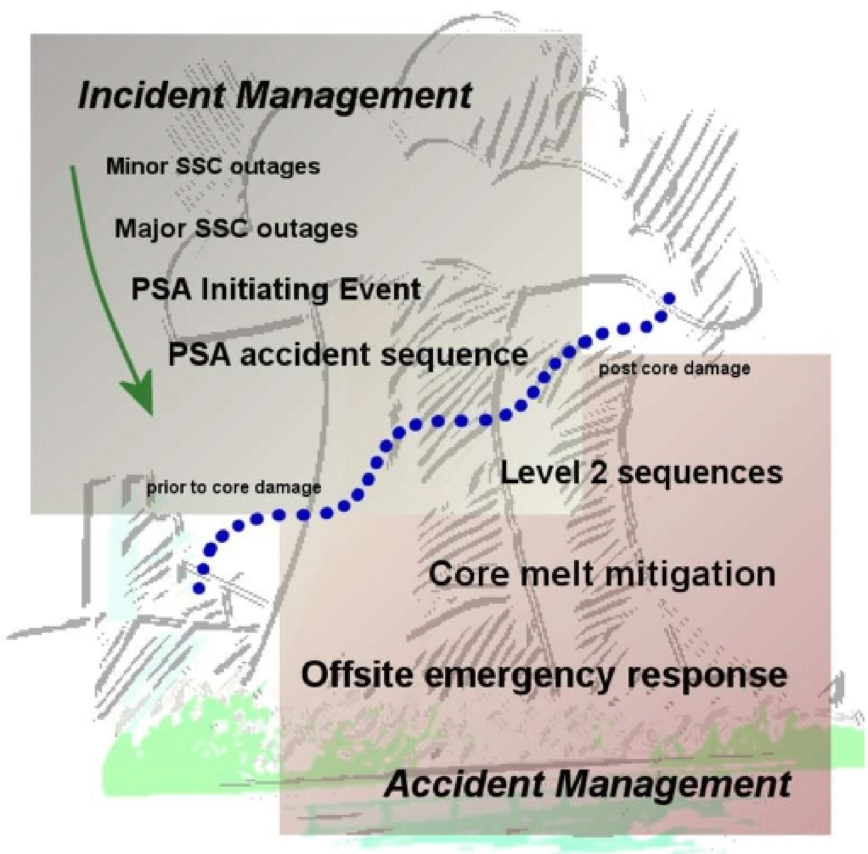

Figure 1. An illustration of the degree of severity between incident and accident management at nuclear power plants.

The remainder of this thesis documents the formal decision making structure we have developed for incident management. Section 2 will provide background information related to both formal decision making and PRA. Then, in Section 3, we will discuss the general methodology development, including utility theory. Section 4 will cover solving the underlying decision model, the supporting analysis modules, and treatment of uncertainty. Section 5 will provide a look at the methodology by way of application to two representative case studies. Finally, we will summarize the primary accomplishments and conclusions in Section 6. Extensive supporting text, including source code to the decision advisor prototype, is provided in the Appendices.

"The difficulty in life is the choice." — *George* Moore

2 Background

2.1 A (Brief) History of Formal Decision Making

Decision making has been part of human activities for numerous years. Lacking predictive information about a future event, decision making historically focused on guessing potential outcomes of one or more realized choices. As decision making evolved, the underlying processes behind the methodology evolved to include probabilistic phenomena. Including the effects of these phenomena – a reflection of a stochastic nature – helped to provide realism in decision making. Consequently, the roots of decision making are tied to fundamental shifts of thought that occurred in the European Renaissance, circa 1600 to 1700, related to the genesis of modern statistical theory. Unfortunately though, many of the important concepts and insights related to decision making did not become widely considered until the middle of the 20th century with publications of von Neumann and Morgenstern's book on game theory (von Neumann and Morgenstern, 1944).

Since the mid 1900s, the implementation of decision making has seen a increasingly strong focus, a focus on integrating decision making theory into mainstream fields such as law, business, and medicine. An important part of the theory behind the implementation is that of probability theory. Like statistics, probability theory has a somewhat recent arrival on the mathematics scene. The early work in this field dates back to the sixteenth through eighteenth centuries and includes Cardano (Cardano, 1663), Bernoulli, and Pascal (Todhunter, 1865). It is worthwhile to note that the Bayesian interpretation of probability theory underlies modern probabilistic safety assessment (PSA).

Building upon the game theory work of von Neumann and Morgenstern, Raiffa and Schlaifer introduced decision analysis concepts such as decision trees (Raiffa, 1968) and how Bayesian mathematics interdict in the application of formal decision theory (Raiffa and Schlaifer, 1961). Couple the fact that computers were becoming available in the 1960s with the emergence of modern decision making theory leads one to appreciate relative rapid rise in popularity of new tools like decision trees. And, as part of these tools, concepts like utility theory were promoted to a new and wide audience of potential decision-makers. In the middle of the 20th century, mathematicians and psychologists explored the utility portion of decision making theory (Edwards, 1954). While the concept of "utilities" is quite old (Laplace, 1825), these theories did not become incorporated into mainstream applications until later where techniques such as the Analytic Hierarchy Process (AHP) and multiattribute utility theory arrived on the scene (Saaty, 1980; Winkler, 1972).

Permeating though all of decision making theory, Bayes theory provides a basis for formulating the decision problem tied to probabilities. Specifically, Bayes provided a technique to process evidence based upon conditional probabilities (Bayes, 1763). Deceptively simple, Bayes' equation states: (Ang and Tang, 1975)

$$p(\theta, x) = p(\theta)\, p(x \mid \theta) / p(x) \tag{1}$$

where $p(\theta, x)$ is the posterior probability distribution; $p(\theta)$ is the prior probability distribution (i.e., what is known about the outcome *prior* to gathering evidence); $p(x \mid \theta)$ is the likelihood of observing a particular outcome or set of evidence (given θ); and $p(x)$ is the unconditional probability of the evidence x (given any θ). We utilize the posterior probability distribution throughout our methodology discussion, specifically on any chance events related to decision making. The uncertainty on these chance events – characterized as *epistemic* – will be discussed in a later section.

The stage for formal decision making has been set by the culmination of statistics, probability, and game theories, but one still has to make decisions based upon imperfect models and a limited state of knowledge. And, these decisions take place on the stage of uncertainty, a place where the dialog, and actors, are to some degree unknown in the next act. It is this fact of making decisions against the backdrop of uncertainty that interjects a pause for reflection, even in presumably straightforward applications of decision making.[†] Four hundred years of science has guided the creation of a formal decision science, a science that helps to guide applications like multiattribute utility theory. Nonetheless, it is important to understand that the science is the foundation, not the entire structure, of formal decision making. The best decision making regimens are ones that mix well defined processes; applicable, informed decision makers; and formal decision-making technologies that assist, but does not replace, human judgment.

2.2 Informal Decision Making at Nuclear Power Plants

A variety of decisions are made every day at every nuclear power plant. Most of these decisions are routine, but, on occasion, significant decisions must be made. Currently, little formal decision making is used in practice, with few exceptions (Weil and Apostolakis, 2001). Nonetheless, informal decision making is used, both by the plant operators and the regulators. It is the intent of this section to provide examples of this application of informal (or ad hoc) decision making practices at nuclear power plants in order to provide a contrast for the formal decision making advocated in this thesis.

[†] Consider the proposition: A game is played by flipping a fair coin until it comes up tails. The total number of flips, *n*, determines the prize, which equals $\$2^n$. For example, if the coin comes up tails on the first toss (which has a probability of 1/2), the prize is \$2, and the game ends. The expected value of the one-toss scenario is: outcome × probability, or $\$2 \times (1/2) = \$1$.   Since there are an infinite number of possible outcomes (*n* heads followed by one tail), each with an expected value of \$1, the expected outcome of the game is an infinite number of dollars. Decision theory states that a risk-neutral person would accept the gamble if the “buy-in” cost is less than the expected return. Since any finite cost is less than infinity, traditional decision theory would suggest playing the game regardless of how large the buy-in cost.

If the initial buy-in were \$2, then of course one would like to play the lottery since you are guaranteed to at least get the buy-in back. But, if the buy-in is a large amount, say \$1,000, one would probably not take the gamble, even though a run of 20 heads and, finally, a tails would return over one million dollars. In other words, one may desire to avoid large losses at the expense of potentially large returns. This concept of “risk” helped to lead to the concept of “utility.” The proposition above is called the ”St. Petersburg Paradox” and is discussed further by Bernstein (1996).

The U.S. Nuclear Regulatory Commission (NRC) regulates nuclear power plant operation through a combination of several regulatory processes. One of these processes, safety oversight, includes activities such as inspection, assessment of performance, evaluation of experience, and other general support activities. As part of these regulatory activities, a variety of risk metrics is utilized. Included in these NRC activities are:

- Significance Determination Process (SDP)
- Generic Issue Resolution
- Risk-Informing Special Treatment Requirements

Let us briefly discuss each in turn.

The SDP is one of the methods that the NRC uses to assist in risk-worth determination of inspection incidents. In general, it is less formal (quantitatively) that other processes. For example, the "Phase 2" SDP uses simplified checklists to estimate an "annualized core damage frequency (CDF)." Decision making then takes place via knowledge of the annualized CDF. Specifically, decision criteria is provide via regions (and associated colors) of interest:

Red	Increase is > 10-4/yr
Yellow	Increase is between 10-5/yr and 10-4/yr
White	Increase is between 10-6/yr and 10-5/yr
Green	Increase is < 10-6/yr

Closely tied to the SDP is the NRC's revised oversight process. This process utilizes metrics such as those discussed above, but then determines a "matrix" of outcomes for measures such as initiating events and mitigating systems. An example of this matrix, called the "action matrix" since the level of regulatory response is keyed to outcomes in the boxes, is shown in Figure 2. But, use of the action matrix raises decision making questions such as consistency between category outcomes. For example, if a plant has two "white" outcomes in a quarter, is this comparable to a "yellow" outcome in another category? Or, is continued impacts (say "white") in one category any better or worse than similar impacts in another category? As we will see later in the discussion of consistency (or sanity) checks for our prototype advisory system, the ability to equate measures and demonstrate consistency is vital.

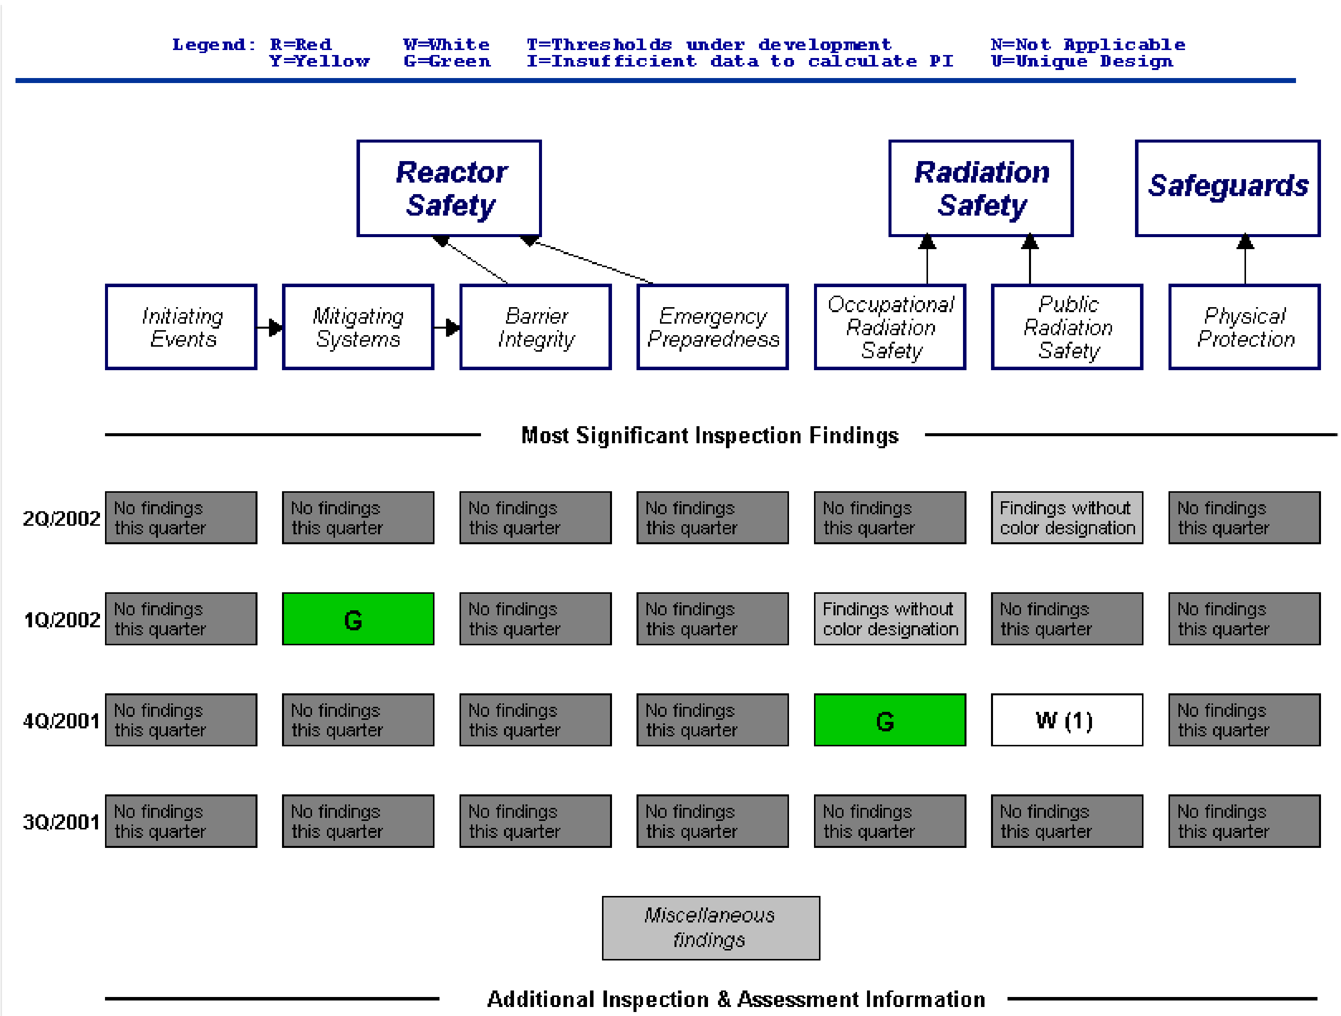

Figure 2. An illustration of the output from the NRC's revised oversight process for nuclear power plants.

The process of Generic Issue Resolution is more formal, quantitatively, than that of SDP. Consideration is taken in calculating both a CDF (specific to various decision alternatives) and the regulatory cost burden associated with decision alternatives. Then, to assess the cost effectiveness of a particular plant alternative, a dollar-to-person-rem averted ratio is generated. Historically, a value of $1,000 per person-rem has been used by the NRC as an upper bound in deciding whether corrective measures may be appropriate. Recently this criteria was changed to $2,000 per person-rem. Additional cost/benefit analysis information is available in NUREG/CR-3568, *A Handbook for Value-Impact Assessment*, and NUREG/CR-4627, Revision 2, *Generic Cost Estimates*.

The Risk-Informing Special Treatment Requirements process is a part of risk-informing 10 CFR part 50 of the U.S. Code of Federal Regulations. Within this process, one proposed option is to make the special treatment requirements (e.g., quality assurance, environmental qualifications, reporting) risk-informed. This proposed modifications would utilize a new definition in 10 CFR 50.2 that depicts which components are "safety-significant." Components that are safety-significant would be within the scope of the requirements. Conversely, components that are deemed to not be safety-significant would be outside the scope unless specifically added by the plant operators or the NRC.

In order to determine the significance of components, the PRA would be used to help determined applicable components using traditional PRA importance measures (Lambert, 1975; Cheok, Parry, and Sherry, 1998). It is desired that the importance measures should

> "be chosen such that results can provide...information on the relative contribution of an SSC (system, structure, or component) to total risk. Examples of importance measures that can accomplish this are the Fussell-Vesely (F-V) importance and the Risk Reduction Worth (RRW) importance. Importance measures should also be used to provide…information on the safety margin available should an SSC fail to function. The Risk Achievement Worth (RAW) importance and the Birnbaum importance are example measures that are suitable for this purpose." (U.S. NRC, 2000)

Proposed decision criteria are based upon these two importance measures, F-V and RAW. If a component exhibits a measure value larger than the target for *either* F-V or RAW, then the component is deemed to be safety-significant. The target importance measure values are:

F-V > 0.005 [for either CDF or large early release frequency (LERF)]
RAW > 2 (for either CDF or LERF)

A central tenant to these risk-informed processes is that the risk model plays a key role to the informal decision making (Brewer and Canady, 1999). Generally, this decision making takes place by use of a risk threshold, for example the values specified on F-V and RAW above. Alternatively, a common risk metric that is used is the change in CDF, or ΔCDF. These metrics then play a part in the larger context of risk-informed applications at current plants. A graphical depiction of the process for risk-informed decision making at nuclear power plants is shown in Figure 3.

Current ad-hoc methods of decision making typically suffer from several flaws. Common problems in these activities include: (1) focusing on a single metric (e.g., importance measures for core damage) as a primary decision-driver, (2) lacking consideration of other decision alternatives outside the initial focus, (3) ignoring decision maker preferences for key attributes, and (4) not using methods such as "sanity checks" to question the validity of decision results. For example, much of the focus of many regulatory programs is on the use of importance measures, where the caveat on their use is typically that the numbers are used to "risk-inform." So, rather than *base* the decision on an importance measure number, the number is applied in an ad-hoc fashion via a subjective process. Note that even formal decision making is subjective, but a formal process forces one to indicate what attribute(s) is important, why is it important, and how much emphasis is paid to the attribute(s) as it relates to decision making.

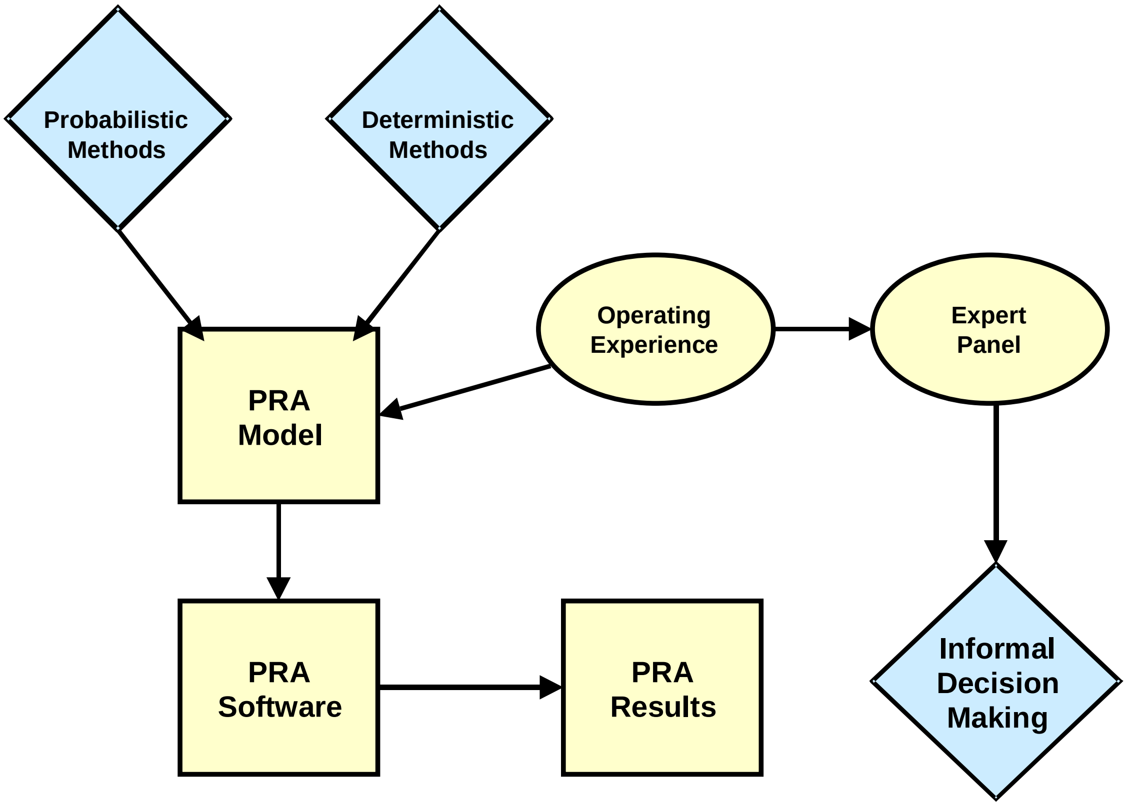

Figure 3. An illustration of the process followed for "risk-informed," informal decision making at many U.S. nuclear power plants.

We will discuss the aspects of formal decision making in more detail in Section 3. But, at this point, we note a couple of identifying features of formal decision making that are absent in many of the informal applications discussed earlier in this section. First, formal decision making integrates a multi-attribute approach (via utility theory) while informal applications typically focus on a single metric type (e.g., in the case of importance measures, the CDF). Second, formal decision making utilizes decision maker preference as an integral part of the analysis. Informal decision making applications generally do not address this issue,[†] and if they do, are somewhat subjective. Third, formal decision making relies on quantitative expectation as the basis for preferential ranking of decision alternatives. Many informal decision making applications focus on “thresholds” (e.g., the colored regions for SDP, \$2,000 per person-rem, F-V < 0.005) which then brings up the question as to the proper decision metric. For example, if the application has a goal that the risk is below a value of *X*, then expected value as a measure may not be desirable since it is a measure of central tendency. Instead, the decision maker may be interested in ensuring that the probability of *exceeding* the risk threshold is low (Smith, Knudsen, and Calley, 1999). In this case, one is focusing on the concept of risk as an upper-bound in relationship to the cut-off criteria.

2.3 An Introduction to Logic-model Based Risk Assessments

Nuclear power plant PRA, like the formal decision analysis models described later in this thesis, is a subjective model. Analysts determine lists of upset conditions (initiating events), the plant response to said upsets (accident sequences), and the performance of specific plant systems (fault trees). Further, as the PRA is decomposed into additional layers of detail, one reaches the lowest level of the PRA, representing individual component behavior (basic events). These individual component modules generally contain either (1) subjective information about a component's likelihood of not performing its intended function, (2) actual failure data, or a combination of (1) and (2). The realm of subjective modeling using probabilistic information falls under the umbrella of Bayesian methods.

At a high level, our Bayesian model is a mixture of deterministic and stochastic (better described by the term "aleatory") modules. For example, both a fault tree and its underlying system success criteria are deterministic. But, because we do not know when a particular component in the system will be inoperable, failures of the component are

---

[†] For example, in the example informal decision making applications that were discussed, questions such as “why is a RAW of 2 significant?” and “how applicable is \$2,000 per person-rem?” are natural and should be addressed with respect to the decision maker’s preference.

represented via an aleatory model. Of course, both these deterministic and aleatory models have parameters associated with them, and each parameter may be uncertain. This second type of uncertainty is classified as epistemic, indicating that our state of knowledge about a portion of the model is incomplete.

To better understand the modeling techniques that make up current PRAs, the major parts of the PRA will be described. In general, a full-scope PRA involves three "levels." The first level contains the logic models (e.g., fault trees and event trees) and probability data representing the outcome of damage to the reactor core. The second level concerns the plant response to the core damage progression (primarily the containment and associated systems). The third level focuses on the off-site consequences resulting from the damaged core and containment. These levels are called Level 1, Level 2, and Level 3, respectively (U.S. NRC 1988). Figure 4 illustrates these three levels and the information that is extracted from each level. Since we are dealing with incident management, we will focus our attentions to only the Level 1 analysis.

Level 1 PRA identifies and quantifies sequences leading to core damage. This process involves identifying significant initiating events, generally those that challenge normal plant operation and that must be successfully dealt with to prevent core damage. Once these initiators are identified, possible plant responses must be determined. The response depends on the different combinations of successes and failures of the systems involved. When the systems have been determined, they must be modeled (usually with fault trees) to identify credible failure modes and unavailabilities. Finally, a Level 1 PRA quantifies the plant's CDF and its associated uncertainty, including at power and reduced power operation. To determine the Level 1 results, initiating event frequencies and equipment failure/unavailability probabilities must be ascertained. This analysis level is the most critical for our decision analysis prototype since the quantitative results are driven by initiators and system unavailability.

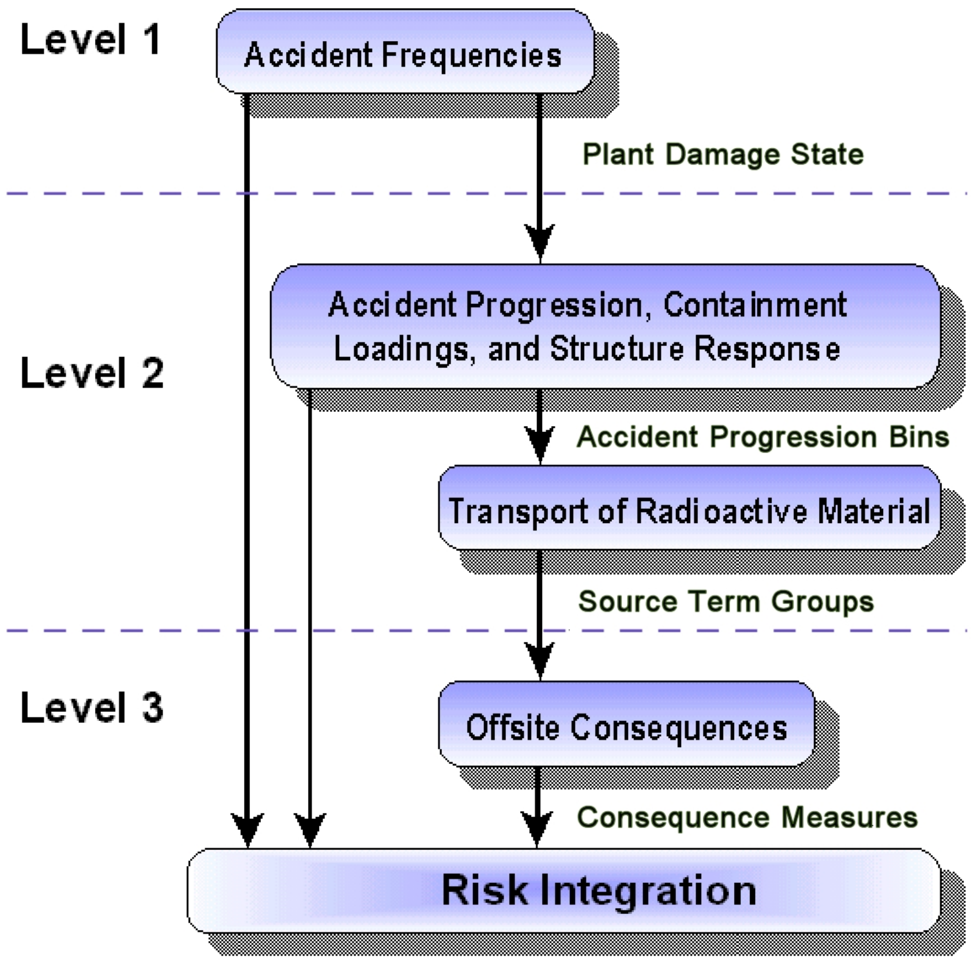

Figure 4. The three PRA analysis "levels."

A Level 2 PRA evaluates and quantifies subsequent material releases from core damage. This analysis involves filtering the Level 1 sequences to a practical number for detailed analysis, typically by grouping Level 1 cut sets into a smaller set of plant damage states. Assessment of containment performance with Level 1 accident sequence analyses is handled much the same as Level 1 analysis by using fault tree models to estimate failure probabilities. A common metric out of the Level 2 models is the LERF.

Level 3 addresses not only Level 1 and 2 issues but evaluates and quantifies resulting consequences to the public and environment as well. Thus, questions such as weather conditions, population levels surrounding the plant site, and dispersion (from containment) characteristics are important in this analysis. Common metrics out of the Level 3 analysis includes early and latent fatalities.

PRA computer tools provide a framework to model traditional Level 1 tasks. For example, event trees can be built to determine accident sequences using initiating events and systems. The individual systems - as named on the event trees - can be modeled using logic fault tree editors. Initiating events and other failure events that comprise each system can be assigned frequencies or probabilities. Minimal cut sets (i.e., a minimally sufficient group of failures that can lead to an undesired outcome) can be generated to quantify fault trees and sequences. The PRA analyst has tools available to perform a variety of different uncertainty analyses, sensitivity analyses, and to calculate importance measures. The capabilities of these PRA tools encompass the following items:

- Initiating events
- Accident sequences (also called sequences or event tree sequences)
- Event trees (also called event tree graphic or event tree logic)
- End states (also called end state partition)
- Systems analysis (also called fault tree analysis)
- Cut set generation (also called cut set solving)
- Uncertainty analysis (also called uncertainty propagation or sampling)
- Importance measures.

These identified PRA areas are considered vital for most traditional PRA analyses. While not discussed here, we do provide additional insights into the structure and utilization of PRA in Appendix A.

As part of the research in defining the decision making prototype, a PRA data transport specification has been developed. This specification utilizes advances in XML (eXtensible Markup Language), a open specification designed to ease data transport over the Internet and between different systems. While the PRA XML definition will facilitate development of our prototype, it is also offered to the PRA community at large in order to provide a “lingua francas” specific to PRA data structures. The full definition of the PRA XML schema is presented in Appendix B. As an illustration of the data structure and nomenclature, we present an example of the XML definition specific to a pump PRA basic event in Table 1.

Table 1. Example of a PRA basic event data structure via the PRA XML specification.

```
<component id="Pump 27-A">
      <title>
      Pump 27-A in the high pressure injection system fails
      </title>
      <failure_information>
            <rate id="operating">
                  <aleatory>
                        Poisson
                  </aleatory>
                  <mean> 1E-4 </mean>
                  <revealed_failure>
                        Yes
                  </revealed_failure>
            </rate>
            <rate id="demand">
                  <aleatory>
                        Binomial
                  </aleatory>
                  <mean> 1E-3 </mean>
                  <revealed_failure>
                        Yes
                  </revealed_failure>
            </rate>
      </failure_information>
</component>
```

2.4 Modules in the Decision Advisor Prototype

After we discuss the theory behind decision modeling and associated analysis heuristics (in Sections 3 and 4), we will develop the specific of our decision advisor prototype. Recall though, that the overall goal of the prototype is to assist nuclear power plant support personnel in response to incidents. Toward that end, we have defined a list of features that should be captured in the prototype software. These features include:

- Elicit the General Context of Incident.
- Construction and Solution of the Primary Decision Model.
- Provide Modeling Information.
- Assist in the Quality Assurance of the Decision Model.
- Provide Results and Sensitivity Metrics.

We have taken these features and used them to construct an information and analysis flow framework. This general framework of the prototype is shown in Figure 5, where a total of five stages are represented. Each stage of the prototype will be turned into modules which, taken together, represent the software behind the prototype. Details on each stage of the prototype will be provided in Section 5.

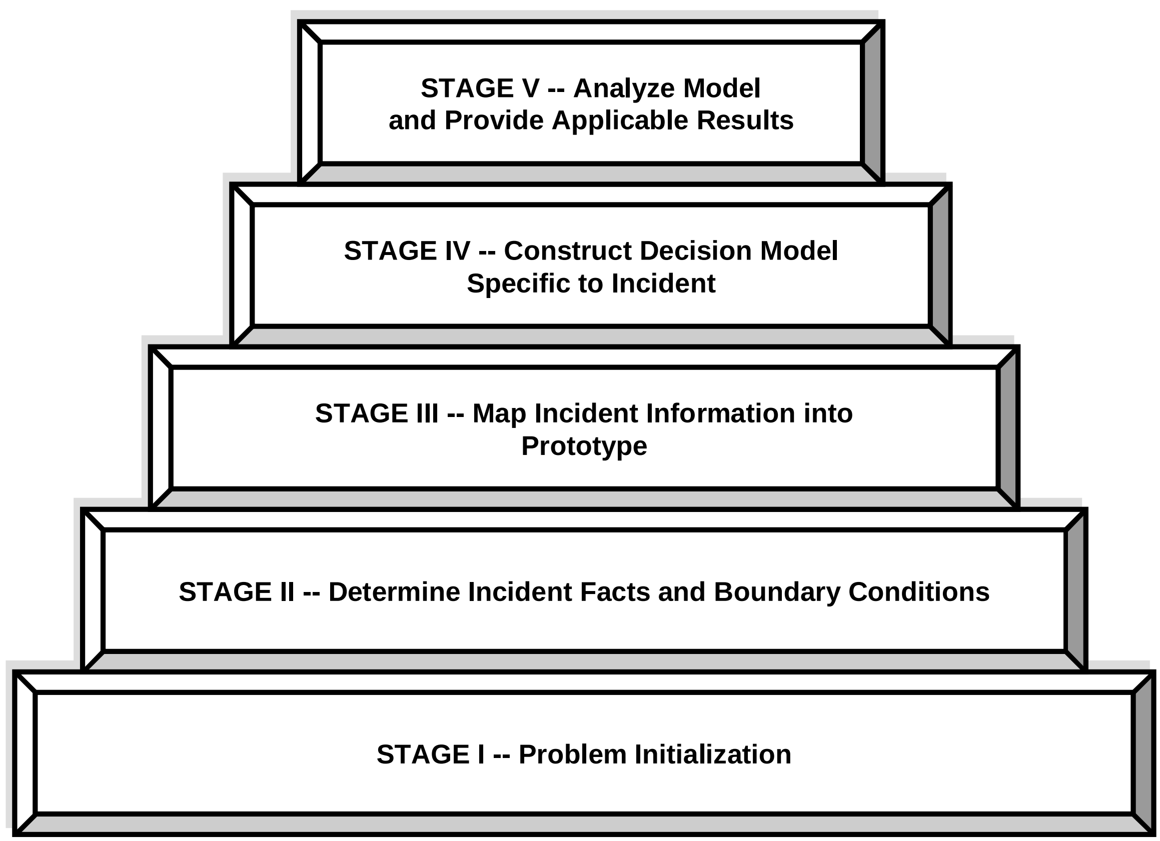

Figure 5. Framework tiers embodied in the decision advisor prototype.

"Life is not numbers…and vice versa." — Anonymous

3 Modeling Methodology for Incident Decision Making

The general methodology behind our incident management framework utilizes formal decision theory with supporting calculation modules such as PRA and disutilities. When faced with an incident, several decision alternatives may be available to the decision maker. It is the goal of the incident advisor prototype to select a preferential decision alternative from the spectrum of available options and provide technical justification for the basis of the decision. For the methodology discussion, we separate the focus into two chapters, modeling of the decision problem (chapter 3) and the subsequent analysis of the model (chapter 4). In this section, we discuss the first aspect, constructing the decision model.

Borrowing from formal decision theory, the general focus for modeling in the incident management framework is through the use of influence diagrams. Briefly, an influence diagram is defined as a directed acyclic graph. This graph utilizes decision nodes (representing the decision alternatives), chance nodes (representing probabilistic events), and outcome nodes (representing the disutilities). The "directed" part of an influence diagram indicates that the various nodes affect one another as dictated by the direction of the arcs (or edges as they are sometimes called). For example, a decision to continue operating the plant may directly affect (or influence) the probability of the core damage chance node. The "acyclic" part of an influence diagram indicates that loops are not allowed. Following a decision, the arcs trace a path from one node to the next, ultimately ending in the outcome node. In this path, we are not allowed to revisit a node. We will discuss the influence diagram (and its related model the decision tree) later in this section. Recognize though that influence diagrams allow the modeler to depict a "big picture" view of the decision problem where this view ultimately is used to present a model of the problem.

Other models that support the influence diagram (for incident management) include modules such as economics, risk, and worker safety. Interacting with these modules is a knowledge base that has been constructed to support both the generation and solution of the incident-specific influence diagram. Combining these two aspects, the influence diagram and supporting modules, with the decision maker preference model provides the overall modeling methodology that is employed in the incident management prototype.

3.1 An Overview of the Decision Problem

We previously classified “incidents” are somewhat benign situation that occur at nuclear power plants. While not serious accidents, they nonetheless demand attention and a technical justification for resolution of the incident. Past ad hoc attempts of decision making did bring some quantitative reasoning to bear, but one of the motivations driving a more formal analysis is the need to better justify decision making. Consequently, we are utilizing the science of formal decision making and mixing that with tools such as the plant’s PRA, an incident knowledge base, and simulation.

Incident management begins with the realization that an incident has occurred which requires attention. Not every upset condition at a nuclear power plant requires formal decision making. Also, many decisions are predetermined due to the regulated environment that governs nuclear power plants. The decision framework we discuss falls within these two decision making extremes. But, the methodology behind the incident decision making always begins with the upset condition as the starting point of the analysis. Then, knowing plant process information related to the incident, the primary decision model can be developed specific to the incident at hand. Augmenting the decision model are modules for decision makers preference and general quantitative supporting calculations. With these three key parts in place, we can then analyze the decision model to determine the preferential decision alternative.

If we focus solely on the decision maker preferences and the supporting calculations, we see several different analysis techniques are required. For example, the preferences module utilizes a model known as a value tree which represents both (1) the decision maker's view on what is important to the decision and (2) how much one attribute is more important than another. These modules that feed into the decision model are illustrated graphically in Figure 6 and will be discuss in detail later in this section.

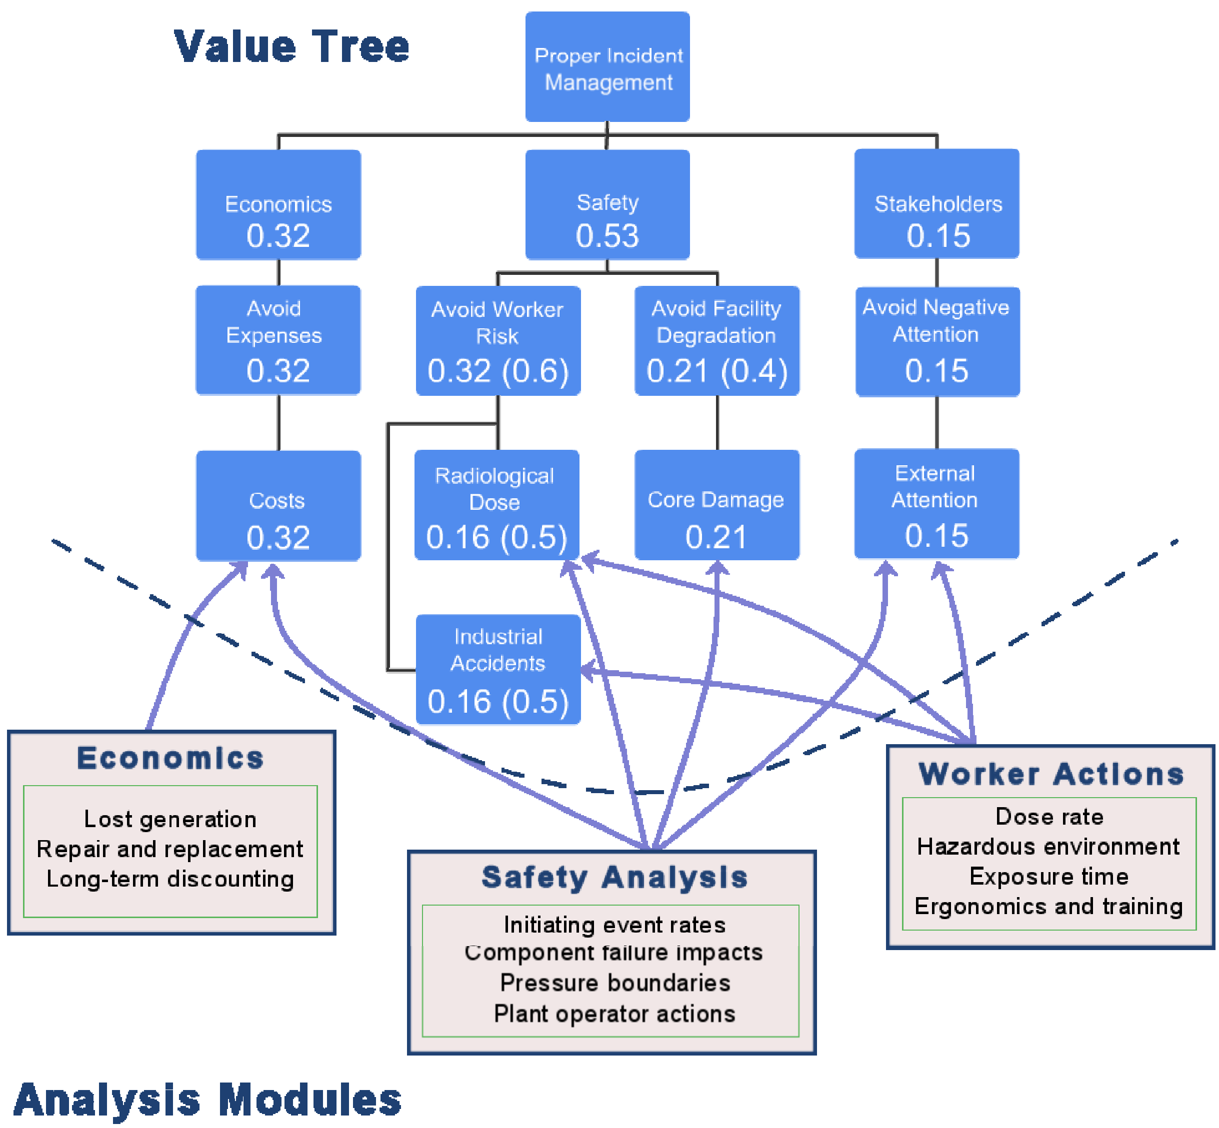

Figure 6. Illustration of the key analysis modules supporting the general decision analysis framework, including key module inputs to the value tree nodes.

We indicated that the supporting modules provide input into the primary decision model. This primary decision model is (generally) an influence diagram. Note though that in place of an influence diagram we could utilize a decision tree since the two are equivalent. That is, if given an influence diagram, one can develop a corresponding decision tree, and vice versa. Prior to discussing the major supporting analysis modules, let us discuss the decision tree model.

In PRA, the two common models for quantification are the fault tree and the event tree. Both are Boolean-based logic models that are typically used to represent system failures or accident sequences, respectively. In formal decision-making approaches, two different models are frequently utilized, namely (1) the influence diagram and (2) the decision tree (Clemen, 1996).

Decision trees and influence diagrams are related models that both provide a structure for depicting decision alternatives and corresponding outcomes. A decision tree is a branching, tree-like structure that represents the results of decision options. Relevant decision options are, by definition, the first branches on the tree. Following each decision branch, we then indicate applicable chance nodes and, at the end of the tree, the ultimate outcome nodes. The decision tree shows, explicitly, all parts of a decision model. This display of the decision model is an asset for small decision models (where "small" equates to less than 6 to 10 nodes) but becomes a hindrance for larger models. Since the decision tree contains all the model detail, many trees can become quite large and cumbersome.

The decision tree defines each decision option and the "sequence" of events following the decision. In a way, the decision tree is similar to a PRA event tree model. In fact, an event tree is a decision tree without decisions. The decision tree sequence consists of the chance nodes in the tree that defines the probability of seeing a certain outcome.

A decision tree contains three important elements. First, the tree begins with the decision bifurcations (decision nodes). These nodes are the first "split" on the decision tree.

There is one branch for each decision alternative. Second, following the decision alternative are the chance nodes representing conditional sequence probabilities. These nodes are the probabilities in the tree. They are generally binary, but may be multi-way splits (note though that the probability should sum to one for each split). Third, the outcome nodes appear on the terminal end of a sequence. These nodes complete the sequence and represent outcomes (e.g., dollars lost, risk, dose, "utility," fatalities). A simple example of a decision tree with these nodes is shown in Figure 7.

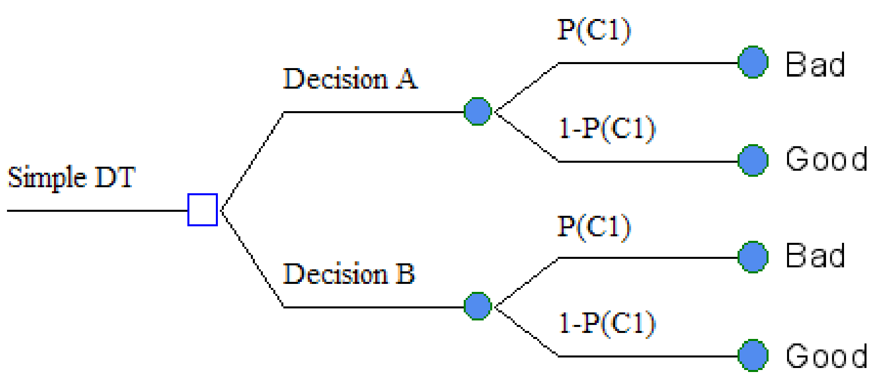

Figure 7. Example of the structure and symbols used in a decision tree.

Given that we have an influence diagram for the decision problem, we can construct a decision tree. Alternatively, one can simply begin the decision modeling via a decision tree. But, for non-trivial decision models, it is recommended that the influence diagram be constructed first since it provides a high-level view of the problem and is easier to manage. Then, the decision tree could be generated programmatically from the influence diagram. In either case, one would utilize the resulting decision model to determine the preferential decision alternative.

For both the influence diagram and decision tree models, the outcome node may be represented by money, or more generally (according to traditional decision theory), a utility function. A utility function is established by the decision maker and represents his or her beliefs and values about a particular attribute of the decision outcome. For example, one may have a utility function for money (the worth of wealth), a different utility function for core damage (the aversion to risky events), and a third utility function for external publicity (the impact of public perception). Once this outcome is defined, the strategy that increases positive outcomes or reduced negative outcomes will be the preferred strategy. Since incident management at nuclear power plants affects multiple, unique attributes that are important to the decision maker, the methodology behind the incident advisor prototype relies on multi-attribute utility theory. We will discuss the implications and limitations of this theory later in this section.

3.2 Characterization of Decision Analysis via Influence Diagrams

As the name implies, an influence diagram is a graphical method to define influences between nodes within the diagram. Here, "influences" can be thought of as dependencies between events, where the dependency may be deterministic or probabilistic. An influence diagram contains nodes and directed arcs between the nodes. The nodes are "variables" specific to the decision problem while arcs identify influences between nodes. The major advantage of the influence diagram is that it shows, at a high level, the "big picture" specific to a decision problem.

An influence diagram contains four important elements. First, a square node indicates a decision nodes (one or more may exist per influence diagram). The decision node indicates that more than one decision options are possible. Second, an oval node indicates a chance node (one or more may exist per influence diagram). Chance nodes indicate a probabilistic event, an event that may occur after the decision is made an acted upon. Third, a diamond node indicates an outcome node (usually only one per influence diagram). Outcome nodes represent the decision makers values about the outcome of a decision. Fourth, the arrows between nodes are arcs and indicate an influence from one node onto another node. The lack of an arc indicates no influence between the two nodes. Using these defined symbols, a simple example of an influence diagram is shown in Figure 8. Within this figure, we show a single decision node (which may contain many decision alternatives), three chance nodes (where node 2 is affected by nodes 1 and 3), and a single outcome node.

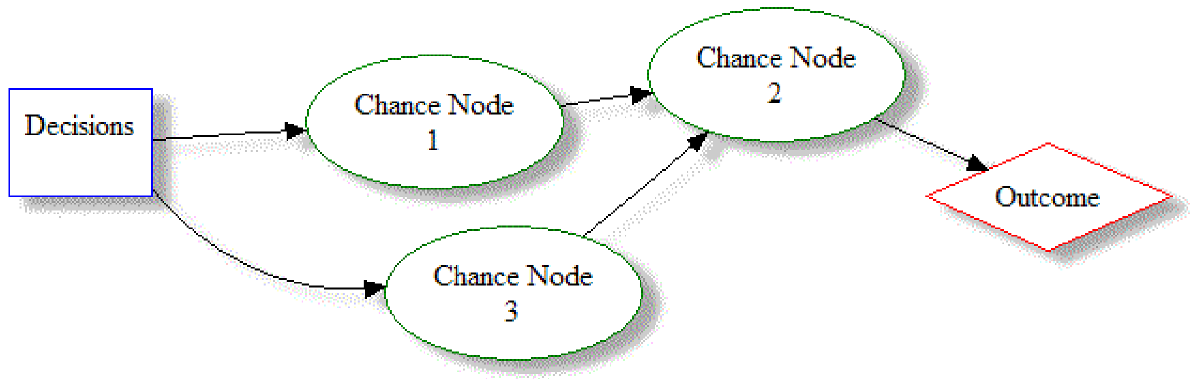

Figure 8. Example of the structure and symbols used in an influence diagram.

One of the guiding principles in constructing the influence diagram is to realize that the model represents the decision maker's state of knowledge prior to making the decision (t = 0). As such, the influence diagram is not a flow chart (e.g., loops are not allowed). Instead, an influence diagram depicts potential outcomes following a decision.

3.2.1 Decision Modeling Within the Prototype

Our incident advisor prototype consists of an expert system that is able to query the user for relevant information applicable to the scenario at hand. Based upon the user input, the advisor system adjusts the analysis case in order to assemble an applicable decision model based upon attribute influences. As part of the model, we will consider items such as uncertainty, the element of time, decision-maker utility, human performance, characterization of decisions, and non-PRA elements (e.g., deterministic characteristics). In general, the influence-based decision model will need to factor in a variety of issues.

One important aspect of the incident advisor is the incorporation of time into the decision scenarios. These time-based, or dynamic effects, can be grouped into two categories: long term effects (such as aging, environmental variations, plant design changes, time until scheduled outages) and short time effects (such as time dependency of physical processes, time to repair inoperable components, operator response times). Many events occurring during an *accident* undoubtedly pertain to the second category, short time effects. But, for an "incident," longer time effects also may be important. For example, if one is considering the option of operating the plant at full power for an extended period of time even though the plant is slightly degraded, then the time to scheduled shutdown may be critical. This operation time may be measured in months to over a year.

It is evident that the incident management influence diagram (and related decision model) will have to deal with a variety of different types (and lengths) of times. Under this assumption, we will classify these times into three natural groups, (1) past occurrence, (2) the present, and (3) future oriented.

The first category, past occurrence, represents passive types of information. What has occurred in the past is (generally) known. Note that during upset conditions such as those described in the PRA, emergency operating procedures, or accident management procedures, details of the immediate past may not be available due to the potential of a large quantity of information in a short duration of time. This "information deluge" has been noted in the literature and could be important during PRA-type situations (Milici et al., 1995). However, for the incident management prototype, past occurrences will most likely represent the boundary conditions for the decision at hand. Consequently, the past occurrence category is important with regard to initially describing the problem for modeling with the incident management prototype. With respect to the influence diagram, one captures the past via boundary condition nodes.

Note though that for the past occurrence category, one should not assume that "past history" implies no uncertainty. Modeling parameters such as repair costs to fix an inoperable component may rely on actual repair data. But, since the repair process exhibits stochastic traits, the prediction of a repair cost is not exact. One will find that uncertainty enters into the decision-making picture in all three of the above categories.

The second category, the present, represents the “here and now” at the time of the decision. During an off-normal incident, the plant process equipment will provide feedback as to the nature of the situation. This process feedback may indicate directional information (e.g., leak rate from the steam generator tube is increasing by X l/hr) in addition to absolute information (e.g., the temperature in the secondary is Y °C). Also, and maybe most importantly, this category (combined with the past occurrence category) represents the decision-makers state of knowledge at the time the decision must be made. From this, one can deduce that the decisions themselves – including the determination of decision alternatives – are encompassed within the time window in this category. Thus, the influence diagram will display the decision nodes, representing the present.

The last category, future oriented, represents the response that occurs as a function of the particular decision that was made in the present category’s time window. It is this category that is most visible in the influence diagram since the majority of the diagram nodes represent events following a decision. The future oriented response generally will include a mixture of hardware and human actions.

Also included in the future oriented category are PRA-related issues such as the potential for core damage conditional upon (a) the boundary conditions at the present and (b) the particular decision to be taken. The use of the Level 1 PRA is an important part of this time window category. In the case where it is very likely that the plant will trip off-line given the present situation, there would not be a large incentive to continue operation to the point that the incident drives plant operation (as opposed to the operators managing plant operation). Other past nuclear power plant management approaches, for example by Milici et al (1995) focused primarily on actions following core damage, or in the realm of severe accident management. In the Milici et al application, their focus was on taking a single event and determining the possible “plant paths” that could follow once the core begins to melt. Consequently, their probabilistic models (given by Bayesian belief networks) represented typical Level 2 PRA models, for example containment pressure, state of water injection systems, and various plant damage states. Also, the decision model focused solely on risk as the performance measure (unlike our application, where plant risk is one of five measures). Our decision model represents the "future" as actions following a decision, where we start from a variety of potential events (all possible initiating events after implementing a decision alternative) that will ultimately wind down into one of two states – core damage or no core damage – as found in our influence diagram outcome nodes. Consequently, our decision model construct is almost a "mirror image" of that found in the accident management literature. In summary, the chance nodes and the outcome node in our decision model represent the future starting from the decision to a predetermined point later in time.

These three time categories, past occurrence, the present, and future oriented, all play important roles in the construction of the decision making model for incidents. The categories also influence the modeling and data the are relevant for the prototype since different items (e.g., the PRA, decisions, and chance node information) may be required at different points in the overall process. Ultimately though, we store information related to these items in the incident advisor knowledge base.

3.2.2 Generalizing the Decision Model

The structure of knowledge base is centered upon storing relationships. In this knowledge base, a "node" represents any possible attribute that may be found in the decision model influence diagram, including decision alternatives, chance nodes, and outcomes. Also, since the nodal information is stored in a relational database, we have the ability to "nest" nodes. While this feature is not found in traditional influenced diagram/decision tree tools, the incident advisor prototype structure allows for the abstraction of high-level decision nodes into sub-nodes. By controlling the node relationships and node types (via its attribute), we can define a hierarchy within an influence diagram. For example, we may have a node identified as "pump system fails to operate" which may appear on the incident-specific influence diagram. But, behind this pump node could be many other sub-nodes representing individual pump failures, common-cause failure mechanisms, operator recovery actions, etc. These sub-nodes would then be utilized to determine the "pump" node failure probability visible in the influence diagram.

Once the knowledge base is populated with plant-specific information related to decision making, we may use heuristics to traverse through the knowledge base. This traversal process then defines the incident-specific influence diagram. An example of code that performs the traversal through the knowledge base is shown in Table 2.

Table 2. Routine for knowledge base traversal for influence diagram construction.

```
<!-- Based upon user input, find starting node -->
Current.Node = Starting.Node
<!-- Find all nodes in the knowledge base that influence CURRENT node -->
Find_Influences(Currrent.Node) = ListBack.Node
<!-- Note that ListBack.Node is an array of Nodes that influence the current node -->
LOOP THROUGH ListBack.Node
        <!—Recursively find nodes that influence current node in ListBack.Node -->
        Find_Influences(ListBack.Node) = NewList.Node
        LOOP THROUGH NewList.Node
                . . .
        END LOOP
END LOOP
<!-- Once all the “backwards” nodes are found, find the “forwards” nodes -->
Find_InfluencedBy(Currrent.Node) = ListForward.Node
<!-- Note that ListForward.Node is an array of Nodes that is influenced BY the current node -->
<!--Add the nodes that have already been found -->
ListForward.Node = ListForward.Node + ListBackward.Node
LOOP THROUGH ListForward.Node
        <!—Recursively find nodes that are influenced BY current node in ListForward.Node -->
        Find_InfluencedBy(ListForward.Node) = NewList.Node
        LOOP THROUGH NewList.Node
                . . .
        END LOOP
END LOOP
```

The influence diagram construction rules listed in Table 2 are illustrated, step-by-step, graphically in Figure 9. In step 1, the initial "entry point" into the knowledge base is known and is used to begin the influence diagram construction. In this example, the incident was the loss of a pressure transducer. In a nuclear power plant, losing a pressure transducer impacts the probability of tripping the plant, hence the "trip|PT" nomenclature (which is read as trip given loss of pressure transducer). In step 2, the routine determines all nodes in the knowledge base that affect or influence the starting "trip|PT" node.

Now that the first level of influences have been found in step 2, then the routine needs to determine all nodes that affect the nodes found in step 2. Consequently, in step 3, the example shows that two nodes affect the "repair" node. In general, this process of determining the "backward" influences is repeated recursively. For our example, the routine requires only three steps to identify all the "backward" influences originating from the starting node. Following the "backward" influence determination, the routine must find the "forward" influences. So, as shown in step 4, we find all nodes that are affected by the "trip|PT" node, where only the probability of tripping the plant is found. This node is now treated as a new "starting" node, where the routine has to repeat the process described in steps 2 and 3. Eventually, after recursively finding the applicable nodes and influences, one is left with the relationships identified in step 6 of Figure 9. These relationships then may form the basis of developing an influence diagram, a decision tree, or the framework for event simulation.

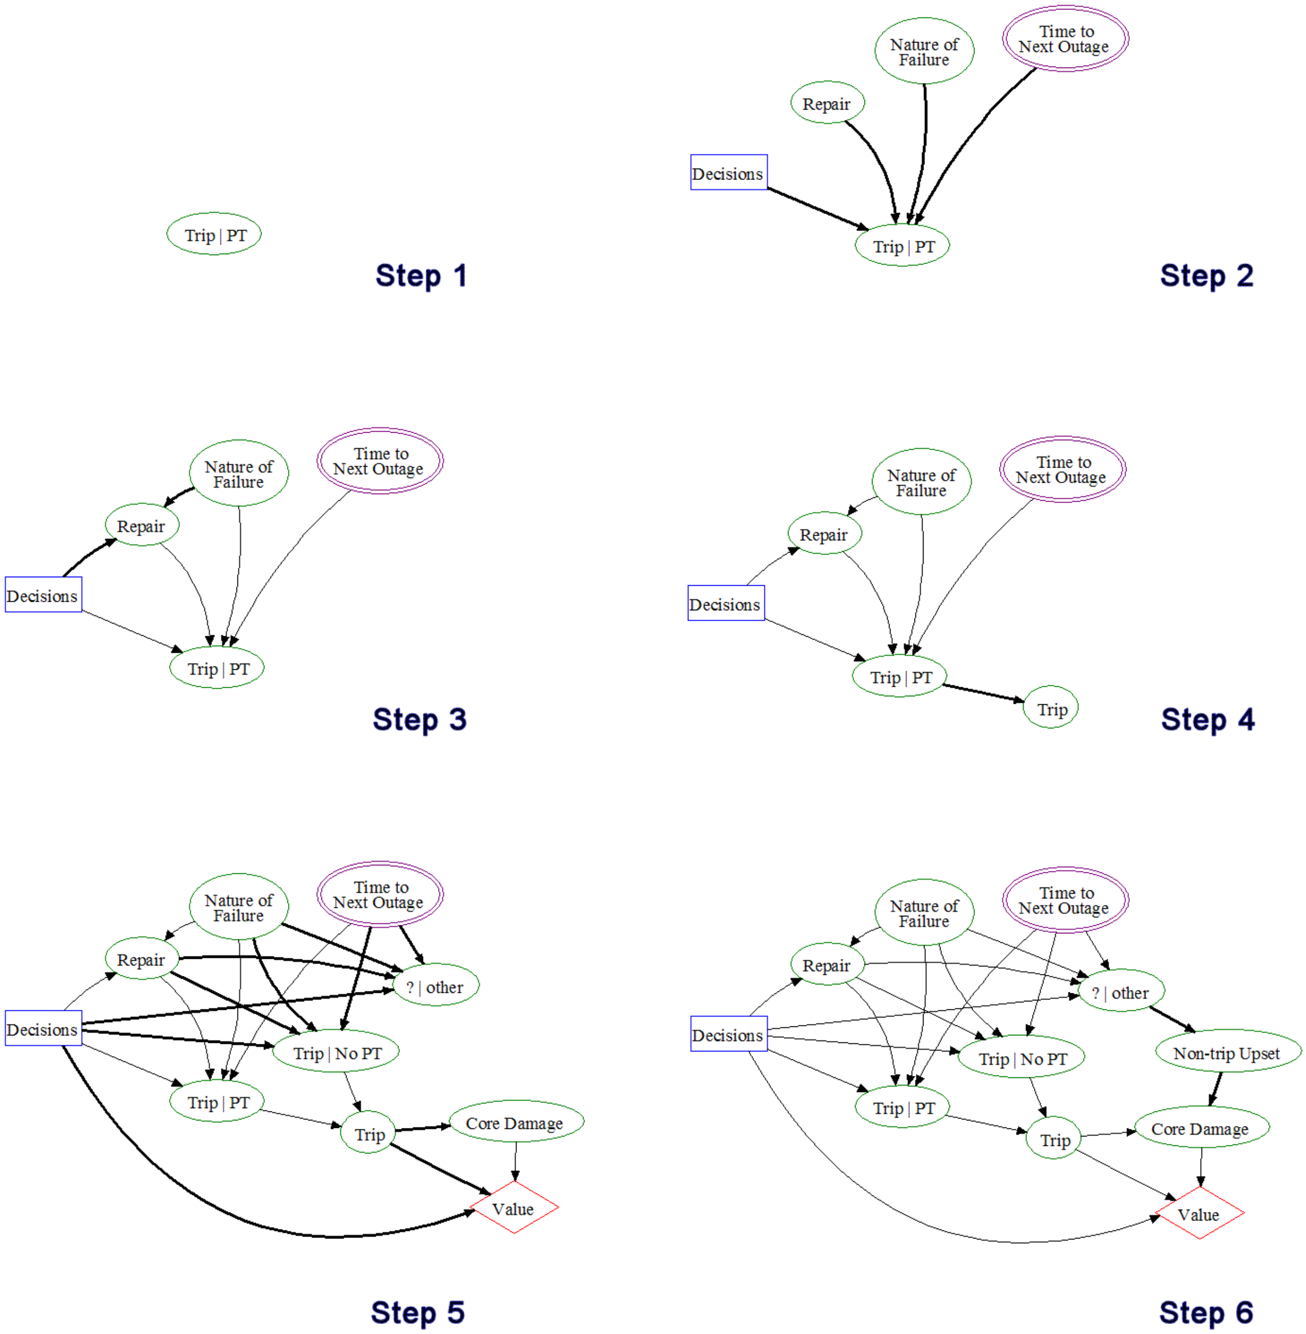

Figure 9. Example of the steps performed by the knowledge base traversal routine for a representative incident case.

Since influence diagrams provide a model representing the “big picture” of the decision process, we can generalize the structure from using the node traversal routine on the incident knowledge base. The results of this generalization are shown in Figure 10. Here we can see that the decision model is made up of six major parts:

- Decision alternatives – these include the options specific to the incident.
- Incident specific elements – these include the possibility for repair, the type of failure mechanisms present, and other unique features related to the incident.
- Boundary conditions – these include the plant state and time until the next outage.
- Plant upsets – these include initiators such as transients/leaks that lead to complications.
- Plant response – these include the plant system response to any upset conditions.
- Outcomes – these include the outcomes of interest to the decision maker.

Now that the general structure of the decision model has been defined via influence diagrams, we need to focus on the other major model supporting the primary decision model, namely that of decision maker preferences.

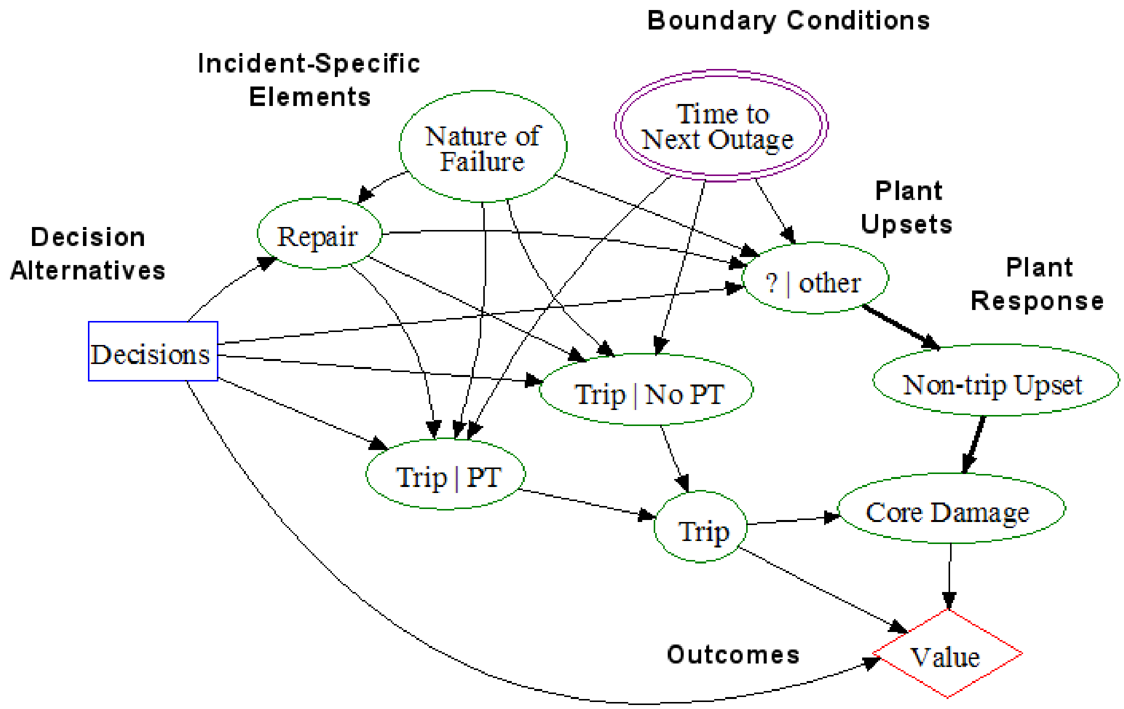

Figure 10. Major elements in a generalized version of an incident-specific influence diagram.

3.3 Decision Maker Preference

In order to describe the key elements of the multiattribute utility decision process, we will describe the process in reverse. First, we will present the theory behind the outcome, expected weighted disutility for all decision alternatives, of the decision process. Second, we will discuss the inputs (e.g., disutility, weights) local to the expected weighted disutility. Third, we will explore the decision maker's risk behavior as it relates to the decision model performance measures.

3.3.1 Application of Utility Theory

Utility theory is a method to translate decision-maker preferences and beliefs, in the context of a known situation, into numeric scores that will ultimately be used to rank decision alternatives. If every decision in incident management had only one attribute of interest, say monetary loss, *and* the decision-maker responded to economic losses in a linear fashion regardless of the magnitude of the loss, we would not need to use utilities. Instead, we could simply use the monetary loss, on a cardinal scale, to rank decisions. But, few decisions, especially those at a nuclear power plant, affect only one attribute of interest. Further, it has been demonstrated via experimentation that decision-makers are not always "risk neutral" (e.g., they do not exhibit a linear response for monetary losses) as the consequences increase (Machina, 1982). Consequently, as part of rational decision making via formal methods, one must rely on the science of utility theory.

Decision researchers subdivide utility theory, and its associated decision making counterpart, into two types of models – "decision making under risk" and "decision making under uncertainty (Knight, 1921; Winkler, 1972)." Decision making under risk embodies situations where one knows the *probability* of any particular state or outcome in the model, but a particular decision-related state is unknown (or exhibits stochastic behavior). An example of this type of model is for the decision on placing a bet upon the roll of a die – assuming a fair die one knows the probability of a specific outcome, but the actual end result of the toss is not known until after the die is cast. Decision making under uncertainty represents the case where the model states are stochastic and the associated probability of any particular state is unknown. An example of this second type of application is for the decision of whether or not to take an umbrella along on a long walk – the possibility of rain is not known exactly and the ultimate outcome is not known until after the decision is made and carried out. In other words, the "under risk" category represents aleatory models with no epistemic uncertainty while the "under uncertainty" category represents aleatory models with epistemic uncertainty. The models that we will encounter for incident management typically fall in the latter case, decision making under uncertainty.

Coupled with the decision model are the criteria one would use when making decisions based upon the model in question. For example, the use of expected value is generally touted as the principal criteria for decision making. But, if we had a decision model with no epistemic uncertainties and the states of the model were know with certainty, one would simply chose the decision with the smallest loss (as measured by disutility). However, since our decision making takes place on the stage of uncertainty, the decision criteria is more complex. Typically, utilization of the expected value approach for decision criteria is applicable only when applied to utilities or applied to outcome metrics when the uncertain outcomes are not too extreme as to cause shifts in preference. But, the use of expected value is not free of complications. We have already discussed the St. Petersburg paradox regarding lotteries with the potential for infinite expectations. Also, the Allais paradox is widely cited as demonstrating a weakness with the use of expected utility (Machina, 1982; Machina, 1987; Conlisk, 1989). Rogers and Fleming note in their work related to nuclear power plant accidents that the "problem with [expected value] is that a single rare event may be catastrophic and financial survival is not assured (Rogers and Fleming, 1998)."

In this thesis, we rely on utility theory to describe our preference in *avoiding negative* outcomes or consequences. Hence, we use the term *disutility* to refer to the numerical value of the outcome of an incident performance measure, where a measure may be described by costs, safety, occupational hazards, and external attention. Since we are focusing our attention on negative outcomes, we ultimately will desire to avoid decisions with (relatively) high values of disutility. As is custom, we bound our disutility between zero and one, where a "zero" disutility implies no or little impact while a disutility of "unity" implies maximal impact.[†]

Since utility theory is a statement of preferential relationships, general rules that dictate the behavior of use for this theory are known. Two basic axioms of utility theory are:

[†] Alternatively, we could have utilized a negatively-valued disutility, where the range would span from 0 (best outcome) to –1 (worst outcome).

Utility Axiom 1 – Preference Ordering. For attribute X, if outcome x_1 is preferred to outcome x_2, then utilities $U(x_1) > U(x_2)$. If outcome x_1 is not preferred to outcome x_2, then $U(x_1) < U(x_2)$. If outcome x_1 is indifferent to outcome x_2, then $U(x_1) = U(x_2)$.

Utility Axiom 2 – Equivalence. If one is indifferent between (1) a certain outcome x_1 or (2) the lottery of outcome x_2 (with probability p) and outcome x_3 (with probability $1 - p$) then utility $U(x_1) = p\, U(x_2) + (1 - p)\, U(x_3)$.

An integral part of formal decision making models is chance or probabilistic events. These events represent uncertain, aleatory outcomes. Dictating the framework for these probabilities are the three axioms of probability theory (Halmos, 1944):

Probability Axiom 1 – Scale Interval. The probability of event X, $P(X)$, should be a number including or between 0 and 1 on a cardinal scale.

Probability Axiom 2 – Certainty. If event X represents a certain event, then $P(X) = 1$.

Probability Axiom 3 – Countable Additivity. If events X_1, X_2, …, X_n are disjoint (i.e., mutually exclusive) then $P(X_1 \cup X_2 \cup \ldots \cup X_n) = P(X_1) + P(X_2) + \ldots + P(X_n)$.

When we couple the axioms of utility theory and the theory behind probabilities with the axioms of coherence (Winkler 1972), we are left with the fundamental basis of formal decision theory. This trio of distinct, but related, axioms then defines formal decision theory. Our set of decision theory axioms is those of coherence, which include:

Coherence Axiom 1 – Ordering of Preferences. For attribute X, outcomes x_1 and x_2 can be ordered by preference, of if they are of equal preference, can be equated to one another. This axiom is effectively the same as Utility Axiom 1.

Coherence Axiom 2 – Transitivity. For attribute X, if outcome x_1 is preferred to outcome x_2, and if outcome x_2 is preferred to outcome x_3, then outcome x_1 is preferred to x_3.

Coherence Axiom 3 – Certainty Equivalent. For attribute X, if outcome x_1 is preferred to outcome x_2, and if outcome x_2 is preferred to outcome x_3, then there exists a probability p_1 such that a lottery of outcome x_1 (with probability p_1) and outcome x_3 (probability $1 - p_1$) is preferred to outcome x_2. Further, there exists a second probability p_2 such that a lottery of outcome x_1 (with probability p_2) and outcome x_3 (probability $1 - p_2$) is not preferred to outcome x_2. Lastly, there exists a third probability p_3 such that a lottery of outcome x_1 (with probability p_3) and outcome x_3 (probability $1 - p_3$) is indifferent to outcome x_2.

Coherence Axiom 4 – Stochastic Dominance. For attribute X, if outcome x_1 is preferred to outcome x_2, and a third outcome x_3 exists, then the lottery of outcome x_1 (with probability p) and outcome x_3 (with probability $1 - p$) is preferred to the lottery of outcome x_2 (with probability p) and outcome x_3 (with probability $1 - p$).

Coherence Axiom 5 – Substitutability. For attribute X, if outcomes x_1 and x_2 are indifferent, then they may be used as surrogates for one another in any decision problem.

Coherence Axiom 6 – *Preference*. For attribute X, if outcome x_1 is preferred to outcome x_2, then the lottery of outcome x_1 (with probability p_1) and outcome x_2 (with probability $1 - p_1$) is preferred to the lottery of outcome x_1 (with probability p_2) and outcome x_2 (with probability $1 - p_2$) only when $p_1 > p_2$.

We illustrate the set of coherence, utility, and probability axioms defining formal decision theory within Figure 11. Within this figure, we note that the resultant decision criteria for decision makers that obey the trio of axioms is that of expected utility.

Figure 12 provides a general overview of the formal decision process described in this paper. In this figure, we show, in a generic fashion, the key portions of the incident management. The output, expected weighted disutility, is quantified as part of the

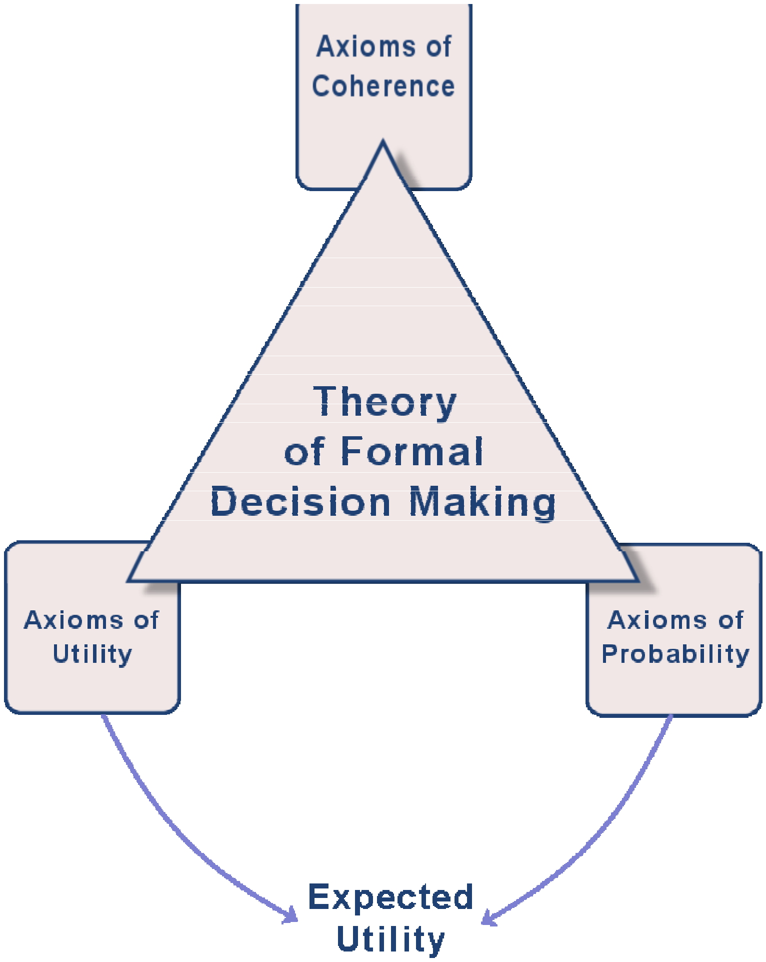

Figure 11. Representation of formal decision making as the amalgamation of coherence, utility, and probability axioms resulting in the decision criteria of expected utility.

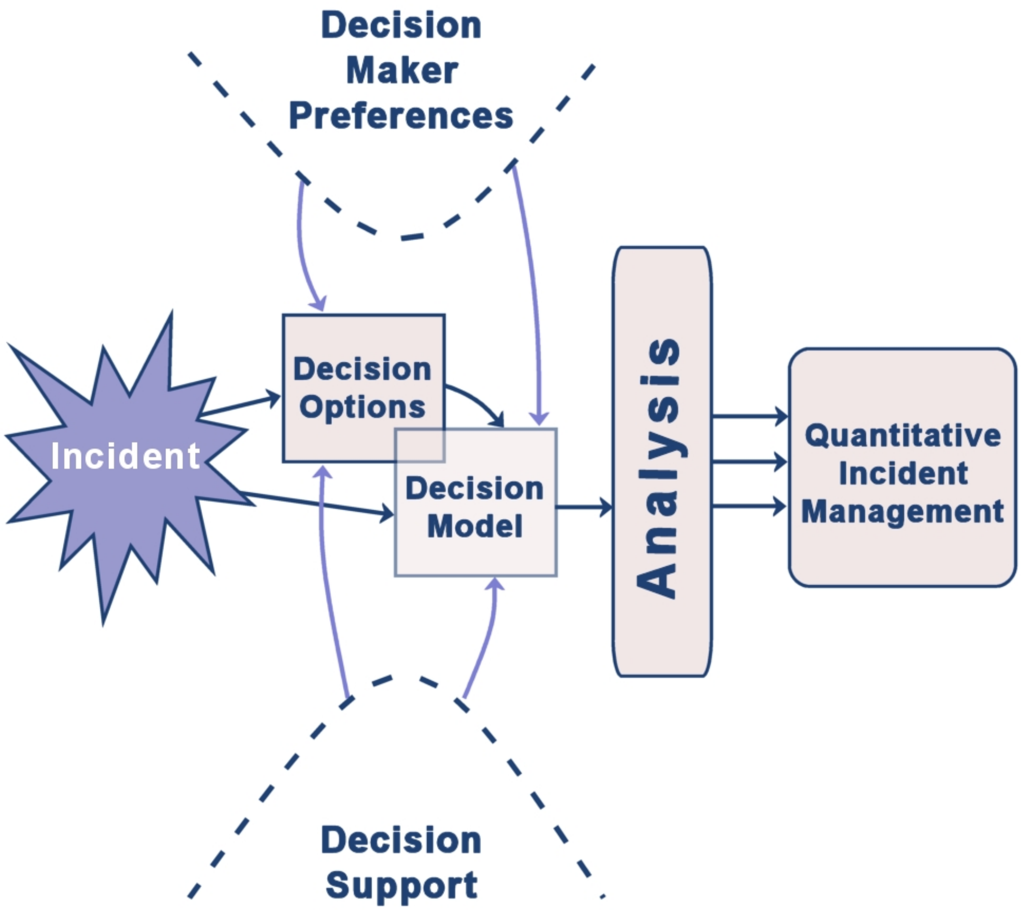

Figure 12. Illustration of the formal decision process beginning with an incident an ending with the quantification of the decision model. The arrows indicate the flow of information specific to that module.

"analysis" box, which is used as one part of quantitative incident management. The other part of incident management is human reasoning and introspection – but the focus in this thesis is on quantitative analysis.

The general form of the weighted disutility, our performance index (also called the "decision rule[†]") is given as a function of two general parameters:

$$PI = f(w_i, u_i) \tag{2}$$

where i indicates a particular performance measure (cost, safety, etc.), w represents a normalized weight, and u is the disutility for the i'th measure. The general expression shown in Equation 2 only serves to describe that the our preference amongst decisions will be a function of performance measures, the weights of those measures, and the disutility associated with decision-specific outcomes.

To develop numerical results, need to describe the interactions between the parameters in Equation (2) in a more illuminating fashion in order to actually make decisions. Thus, we turn to a slightly more interesting version of the performance index (Keeney and Raiffa, 1993; Winkler, 1972) for the case of three performance measures X, Y, and Z:

$$PI = w_1\, u(x) + w_2\, u(y) + w_3\, u(z) + w_4\, u(x,y) + w_5 u(x,z) + w_6 u(y,z) + w_7 u(x,y,z) \tag{3}$$

where w_1 is the weight for attribute X; w_2 is the weight for attribute Y; w_3 is the weight for attribute Z, w_4 is the weight for the joint outcome of X and Y; w_5 is the weight for the joint outcome of X and Z; w_6 is the weight for the joint outcome of Y and Z; w_7 is the weight for the joint outcome of X, Y, and Z; $u(x)$ is the marginal (or unconditional) disutility of attribute X; $u(y)$ is the marginal disutility of attribute Y; $u(z)$ is the marginal disutility of attribute Z; $u(x, y)$ is the joint disutility for X and Y; $u(x,z)$ is the joint disutility for X and Z; $u(y, z)$ is the joint disutility for Y and Z; and $u(x, y, z)$ is the joint disutility for X, Y, and Z. In this paper, we use the notation that X is a particular attribute, x is a specific level of that attribute, and $u(x)$ is the disutility associated with a specific level of attribute X.

[†] The decision rule, or decision criteria, really is manifested by determining the *expected value* of the weighted disutility, which we then uses to rank decisions (where low values are preferred).

During the determination of attributes for incident management at nuclear power plants, the decision maker decided on a total of five performance measures: economics, radiological dose, industrial accidents, safety, and stakeholder attention. Consequently, to derive the full equation for the performance index would result in a functional form with 31 terms, each with their own weight. In decision science, it is typically argued that the relationships described in Equation 3 should be simplified for a variety of reasons, including:

- The decision results are used to advise (e.g., risk-*informed* applications) rather than decide (e.g., risk-*based* applications), thereby allowing errors or discrepancies in the process to be discovered and overcome
- The decision attributes are *fundamental*, thereby prohibiting interaction between the attributes.

Unfortunately, none of the above points provide technical justification for simplifying Equation 3. Ideally, one would strive to determine the tradeoff between simplicity against predictability. However, the science behind model selection is beyond the scope addressed in this thesis. Consequently, we will not enter into the discussion of issues like predictive accuracy (Forster, 2001) or “knowability” (Casti,1990). Instead, we will point out that we can greatly simplify the performance index function if the performance measures themselves are *additive independent*, where additive independence also implies *utility independence*. (Keeney and Raiffa, 1993) While the additive model may be considered to be an approximation (Clemen, 1996), in many cases the model is adequate.

For the special case of additive independence, when we have three performance measures X, Y, and Z, the performance index becomes:

$$PI = w_1\ u(x) + w_2\ u(y) + w_3\ u(z) \tag{4}$$

or, in general, we can simply sum the weighted disutilities for each performance measure.

For our case with incident management, we would then have:

$$\begin{aligned} PI = {} & w_{economics}\ u(economics) + w_{dose}\ u(dose) + w_{accidents}\ u(acc.) \\ & + w_{safety}\ u(safety) + w_{stakeholders}\ u(stakeholders) \\ = {} & PI_{economics} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{stakeholders} \end{aligned} \tag{5}$$

Note that it has been pointed out in the literature that for additive independence to apply "there cannot be any interactions among the performance (attribute) measures" (Weil and Apostolakis, 2001). Interactions between the attributes *are* allowed, depending on the type of interaction. For example, probabilistic dependency between the attributes is not only permissible, but to be expected. A decision alternative that results in an accident is likely to have high costs, high safety impacts, and a high level of external attention. But, there are other dependencies that one must consider when utilizing the expression in Equation 5. In general, there are four kinds of independence that play a role in DM:

- Probabilistic independence
- Preferential independence
- Utility independence
- Additive independence

Probabilistic Independence. This type of independence is the standard probability definition of independence between two random variables that you would find in any probability theory textbook. For example, for events A and B, if they display probabilistic independence, then $P(A \cap B) = P(A) \cdot P(B)$. In our case, if two performance measures (X and Y) are probabilistically independent, we are allowed to describe their expectation as:

$$E[u(x1, y0) \cap u(x0, y1)] = E[u(x1, y0)] \cdot E[u(x0, y1)] \tag{6}$$

where $x0$, $x1$ are specific levels of the X performance measure and $y0$, $y1$ are specific levels of the Y attribute. Note that probabilistic independence is an issue related to interactions at a probability level, and as such, impacts the determination of metrics such as expectation and variance. While a subtle point in the context of decision making, the aspect of probabilities being *separate* from preference (and disutility) is critical to both the development of appropriate disutilities, attribute weights, and the quantification of weighted disutility. Probabilistic dependence does not interject itself into the topic of model selection for weighted disutility; instead we have other independence (conditional, preferential, utility, and additive) conditions to guide the form for weighted disutility.

Preferential Independence. This type of independence represents static preference of two utilities in the presence of a third attribute. Attributes X and Y are preferentially independent if $u(x, y)$ given $u(z0)$ does not depend on the particular level $z0$ where X, Y, and Z are attributes and $z0$ is a specific outcome for a attribute Z. For example, assume X is cost, Y is loss of property, and Z is availability of donuts in the on-site cafeteria. For incident management at nuclear power plants, a decision maker preference for cost versus property loss will be indifferent to the availability of donuts, or:

$$u(\{x \mid z\}, \{y \mid z\}) = u(x, y) \tag{7}$$

While this type of independence may seem uninteresting (since we obviously do not care about donuts in the context of incident management), this property allows us to determine that many attributes may be irrelevant, thereby simplifying the decision process. In other words, when making complex decisions, we do not need to have a large set of interests (in general, the set of interests is infinitely large) – we only need to focus on attributes that sway decision alternative outcomes.

Utility Independence. This type of independence is used to explore potential preference interactions between utility attributes. For example, X is utility independent of Y when conditional preference for utility X (given y) does not depend on y. In other words, a particular attribute outcome does not affect the preference amongst attributes, even if the outcome is extremely negative or positive. This restriction implies that, given a range of values for one attribute, a second attribute's utility (or disutility) does not change. Note though that if X is utility independent of Y, the converse is not necessarily true – when the converse is true, this condition is denoted as mutually utility independent. When combining mutually utility independent attributes, Equation 3 is applicable.

With our decision-makers, we determined the five performance measures relevant to the incident management context. We later ran an experiment to test the notion of utility independence, specifically within the context of safety. We deliberately selected two measures that, although not utilized in the current framework of our weighted disutility, could possibly demonstrate an extreme case where attributes *were not* utility independent. First, we denoted one performance measure as the strength of containment,[†] where we had three scales for containment performance: (1) robust, (2) moderate, and (3) weak. Second, we set the other performance measure as core damage probability, and fixed this attribute at one of two levels: unlikely (a probability of 1E-7) or likely (a probability of 0.2). We then elicited the containment performance disutility first assuming that core damage was unlikely. It was then communicated to the decision-makers that the disutility-determination process was to be repeated, but they were to assume that core damage was likely. The results of this experiment, as demonstrated by the two conditional disutilities, are shown in Figure 13.

The two disutility curves for containment performance indicate that our decision attributes of core damage and containment performance are not *exactly* utility independent since the two curves do not coincide. But, we were surprised that the two curves do not deviate by a substantial amount (the maximum deviation is about 31%, which is within the variability of the weight determination process -- see Appendix C for details on this point). Since containment performance is critical if core damage is likely, we would have thought it natural that preference on performance would be correlated

[†] The nuclear power plant containment serves as a defense barrier to prevent release of radioactivity in the event of a severe core melt accident.

with likelihood of a core damage accident. Instead, preference for containment performance was better described as being of fundamental importance regardless of accident likelihood. Hence, as an approximation, we could model these two performance measures as being utility independent. And, of all the performance measures that we have previously discussed for incident management, it was thought that these two, core damage and containment performance, were candidates for not being utility independent. Thus, we may be justified in assuming that the remaining performance measures are (approximately) utility independent.

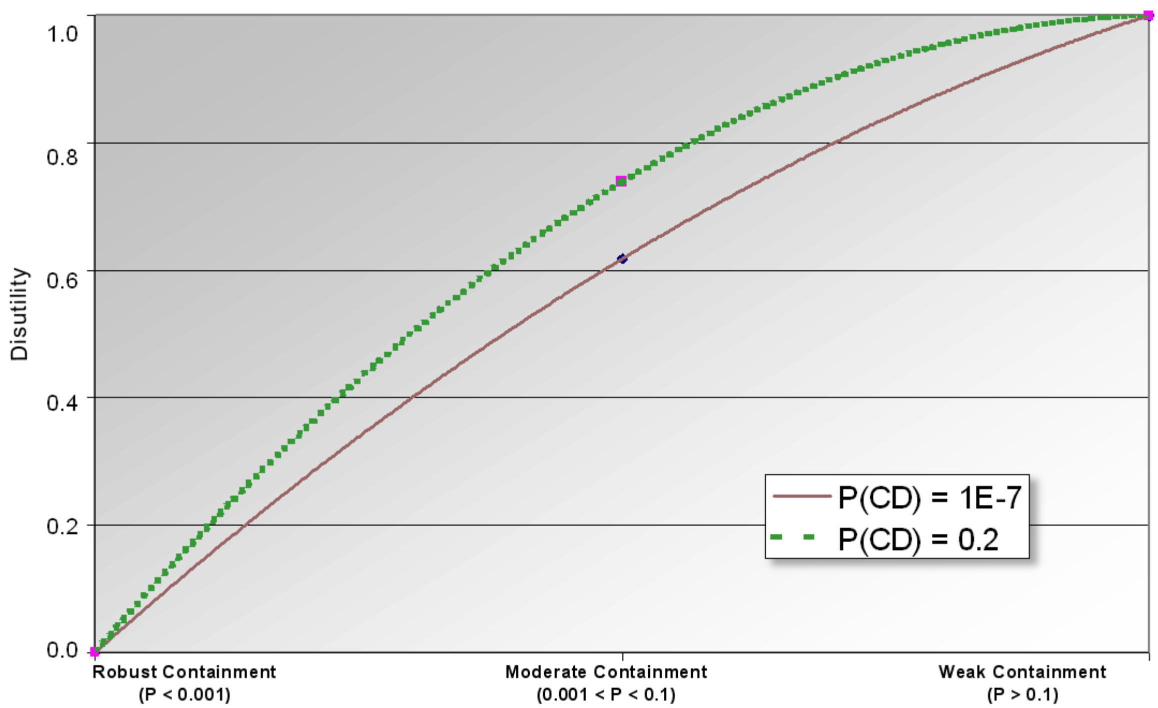

Figure 13. Two performance measure experiment test results for utility independence.

Additive Independence. This type of independence represents preference amongst utility lotteries (i.e., alternatives represented by known utility outcomes and associated probabilities). If attributes X and Y are additive independent, then the decision-maker will be indifferent between the two lotteries shown in Figure 14.

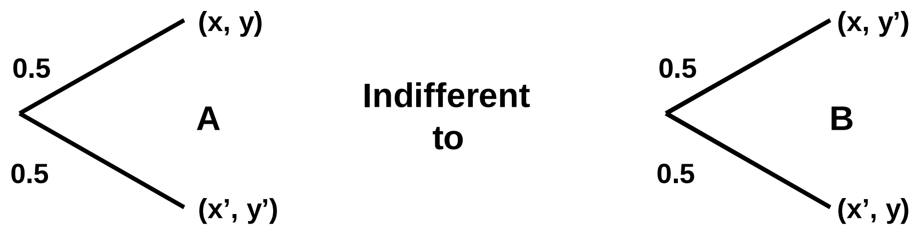

Figure 14. Lotteries used to test utility functions for additive independence.

For example, if X is cost and Y is safety, then outcome x may represent the lowest cost, outcome x’ may represent the maximal cost, outcome y may represent the safest condition, and outcome y’ may represent the most risky condition. If X and Y are additive independent, then the decision-maker must be indifferent between (1) a 50-50 lottery where the “payoff” of lottery A is lowest cost and lowest risk *or* the highest cost and highest risk and (2) a 50-50 lottery where the “payoff” of lottery B is a combination of low cost/high risk *or* high cost/low risk.

Additive independence is the strongest of the four dependence types and leads one to an additive utility function. In other words, if attributes X, Y, and Z are additive independent then we could utilize Equation (3) by observing that the joint term weights are zero, or $w_4 = w_5 = w_6 = w_7 = 0$. Note that additive independence implies mutual utility independence, but the converse in not necessarily true. Decision makers strive to make use of attributes that are additive independent due to the fact that they are then allowed to use the function form of weighted utility described in Equation 5. Alternatively, many practitioners simply invoke the principle of Occam’s razor (one should not make more assumptions, or complications, than the minimum needed). While Occam's razor allows us to "drop" variables or model constructs that are not really needed to explain a particular domain outcome, one runs the risk of oversimplifying a complicated methodology, thereby jeopardizing its predictive capabilities. But, in lieu of attempting to utilize an overly complicated decision model, Clemen (1996) suggests

"…in extremely complicated situations with many attributes, the additive model may be a useful rough-cut approximation."

As a side issue, we assume that the decision-maker preferences – and corresponding disutility – do not change over time. Researchers like Keeney and Raiffa do propose utilities for time streams of unequal length, where a "family of conditional utilities" may arise (Keeny and Raiffa, 1993). But, the contexts of these situations were assumed to be over long periods, where examples included pension, career, and savings decisions. While cases have been considered where preference varies over time (Phelps and Pollak, 1968), it is not clear that such implications would apply to incident management at nuclear power plants (where it is desirable to have less subjectivity, hence the introduction of formal decision methods) since the time window is short. This restriction does not imply that if the context of the decision problem changed that preferences and disutilities would not change. For some non-incident situations (say in the case of a severe accident), it may be easy to postulate that the decision-maker's focus would shift from a balanced approach of cost, safety, worker health, etc., to one predominately centered upon the safety of the public around the plant. Hence, it is important to once again point out that the issues and experimental data presented in this paper are solely for the context of incident management at nuclear power plants.

3.3.2 Inputs to Expected Disutility for Incident Management

With the discussion from the previous section, consistency checks on the decision attributes (which will be presented), and discussion with our decision-maker, we were led to the assumption of using a linear, additive weighted disutility for the performance index:

$$\begin{aligned} PI &= w_{cost}\, u(cost) + w_{dose}\, u(dose) + w_{accidents}\, u(acc.) \\ &\quad + w_{safety}\, u(safety) + w_{attention}\, u(att.) \\ &= PI_{cost} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{attention} \end{aligned} \tag{8}$$

As part of this function, we will need to determine both performance measure weights (*w*) and associated disutilities (*u*). Once we have this information, we will be able to question the "linear additive" assumption. For example, we may develop multidimensional indifference curves that illustrate tradeoffs from one attribute to another. These so called "sanity checks" will allow us to ensure consistency between (presumed) independent attributes. If inconsistencies cannot be resolved, it may be necessary to rectify independence issues by returning to a more complicated functional form of the performance index.

We have already indicated that the applicable decision performance measures to incident management were costs, safety, occupational hazards, and stakeholder attention. Questions that arise are "how are these measures derived?" and "how do we obtain the weights for each measure?" The short answer to each question is (1) through deliberation with the decision-maker we determined applicable performance measures and (2) through the application of AHP we derived measure weights. Since decision making is a deliberative process, the outcomes of the process may evolve as new insights or discussions are presented. For example, during our workshops with our decision makers, the structure and node in the value tree were modified through three revisions. We would like to point out important modifications to the value tree that resulted over the course of these revisions.

> First, the economics category had two performance measures that were ultimately combined into a single category. We originally identified two types of costs, short-term costs and long-term costs. The rational behind the separate categories was that for some incident related decisions, the decision-maker would incur up-front (realized) costs. But, depending on future (uncertain) events, an incident could develop into a worse condition and could potentially result in much higher costs (for example, in the case of a accident or unchecked degradation that damages expensive equipment). In other words, we were addressing both deterministic costs (i.e., up front costs) and probabilistic costs (i.e., potential future costs). But, a dollar now should be worth a dollar (discounted

appropriately) later. Keeny and von Winterfeldt acknowledge this point by noting that "future costs are discounted to a present value equivalent" (Keeney and von Winterfeldt, 1994). Consequently, in later revisions to our value tree, the two cost types were combined into a single "cost" category.

Second, like the economics category, the stakeholder portion of the value tree originally had two attributes, the regulator and the media. These two categories were ultimately combined into a single category, stakeholders. It was felt by the decision-maker that negative attention from outside, the regulator, the media, or the public, was undesirable and, in fact, may be highly correlated. It was thought that, for incident management, the regulator would be the first party involved (if any) since these events do not typically result in "high profile" situations. Then, if the regulator is concerned, the media attention may increase. Consequently, it was decided that a single attribute, stakeholders, was justified and should focus mainly on interactions with the nuclear regulator. Note though that the delineation and weight placed on this attribute may vary from country to country.

Third, the "core damage" performance measure under the safety objective was modified more than once. Originally, the attribute was core damage *frequency*, but was changed after the first revision to core damage *probability* to reflect the fact that different incidents may span vastly different time periods. In general, the longer the time period, the higher the probability of core damage. But, like the probabilistic problem we faced on the "cost" issue, it is problematic to include probabilities directly in the value tree. This "probability-in-the-value-tree problem" is not unique to low probability, high consequence situations like running a nuclear power plant, as evidenced by the difficulty experienced by Wright and Goodwin in their "new job" problem (Wright and Goodwin, 1999). Ultimately though, as noted by others, the value tree should (as its name indicates) reflect the decision-makers values (Keeney and Raiffa, 1993; Clemen, 1996). Consequently, one should strive to avoid introducing probabilistic information (however well intentioned!) directly in the value tree and corresponding decision attributes. All probabilistic interactions should take place when quantifying the performance index, for example, when determining its expected value.

To summarize the value tree evolution, we removed probabilistic aspects (potential costs, probability of core damage) in order to focus on *observable* events. While not explicitly stated in much of the current decision analysis literature and texts, we believe that construction of a decision-maker's value tree should have attributes that are measurable. Second, we identified, through deliberation, that media perception and regulatory interest *were not* preferentially independent. This observation allowed us to reduce the number of attributes (by subsuming one), yielding a slightly less complex value tree and weighted disutility function. Note that deliberation played a critical role both in shaping the structure of the tree and in determination of the performance measure weights. As part of this deliberation, we performed "sanity" checks of the nominal results with the intent that the decision makers could see the overall impact from their discussions. In one case, it was these sanity checks that caused a group of three decision makers to adopt the position of a fourth (sole) decision maker, mainly due to the fact that the results appeared to be more reasonable for the lone decision maker's value tree.

With the functional form of the performance index now set and the value tree finalized, we may proceed to determine (1) performance measure weights and (2) performance measure disutilities.

3.3.3 Performance Measure Weights

The performance measure weights were determined by utilizing AHP for pair-wise comparisons at each level in the value tree (Saaty, 1980). The value tree, with the resulting weights, is shown in Figure 15. During the AHP process, the preference comparisons were always made with respect to the next highest level. For example, when determining the weights on economics, safety, and stakeholders, the context was with respect to "proper incident management." Then, following AHP, deliberation

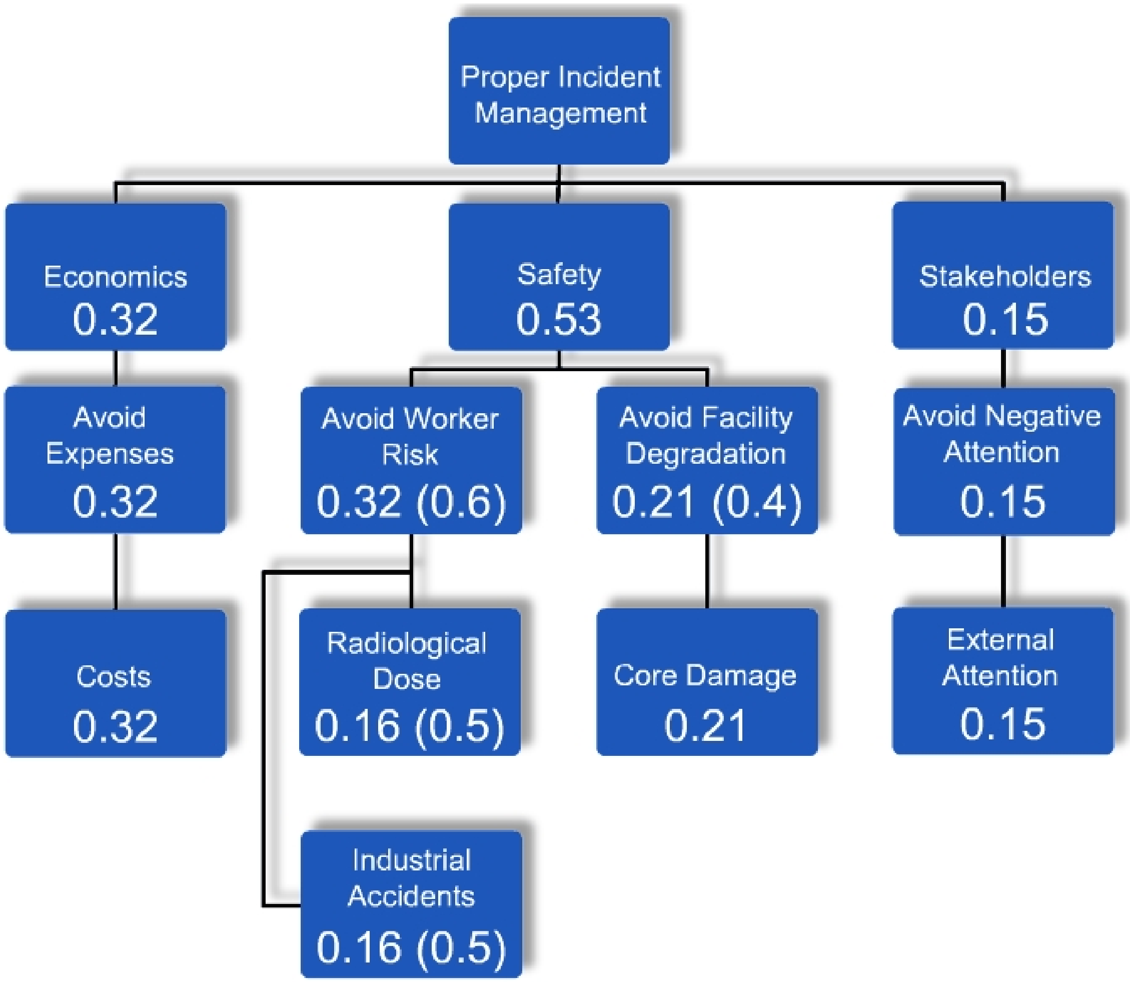

Figure 15. The value tree with performance measure weights determined from the AHP/deliberation process.

was used to modify the weights to better reflect the decision-maker preferences. The results of this AHP-deliberation process indicated that the decision-maker values the utility of economics over stakeholders by about a factor of two (0.32/0.15), while that of safety is not quite a factor of two over economics (0.53/0.32).

As an aside, we note that a related study was found after our analysis was complete that focused on decision making *following* a nuclear accident (Lindstedt et al, 2001). While the context here is somewhat different that ours, the researchers and decision makers nonetheless developed a value tree that, at the first level, was quite similar to ours. Their determined weights were cost = 0.26, health = 0.43, and sociological = 0.32. Slightly less emphasis was placed on cost and health, but more on external impacts. This bias in weights toward external impacts may be due to the fact that here the decision makers were not the operators/owners of the nuclear power plant, but instead were members of the public (e.g., farmers, dairymen).

To determine both the structure and weights associated with each node, we held a workshop with the primary decision makers. A list of the activities during the workshop is shown in Table 3. The deliberation phase of the project proved useful for revising the performance measure weights. Originally, the decision makers were split into two camps where three of four provided weights that were similar while the fourth decision maker had substantially different weights. During the deliberation, we presented consistency checks for the preliminary results that lead the group of three to revise their weights closer to those of the single decision maker. If we had taken an average of the four sets of weights, the results would have been much closer to the group of three weights rather than those shown in Figure 15. We found that the combination of consistency results combined with deliberation yields a useful procedure for the ultimate determination of decision maker preferences.

Table 3. Activities performed during the preference elicitation workshop.

Topic	Purpose
Introduction to decision methodology	Ensure that decision makers understand the context of the preference elicitation
Discussion of AHP	Discuss implementation of the AHP
Development of the value tree	Finalize the value tree including the general structure and nodal weights
Development of the utility functions	Elicit preference information for each of the value tree performance measures
Presentation of weights for final deliberation	Provide a forum for deliberation of the value tree weights (note that deliberation of the utility functions was held in a later workshop)

3.3.4 Performance Measure Disutilities

Following the structure and weight determination of the value tree, the disutility functions need to be constructed. First though, let us define the types of scales of interest for our disutility functions. In general, within the decision sciences, there are three types of scales that are utilized, cardinal, ordinal, and nominal.

Cardinal scales represent specific outcomes (e.g., in our notation for attribute X, values associated with x and x') where the *difference* between outcome values has meaning to the decision maker. For example, losing one million euro versus losing two million euro is a quantifiable difference and will impact profitability. Within the cardinal scale category, we can subdivide the scales further into interval or ratio types. If the location on the cardinal scale is arbitrary, the scale is called an interval scale, otherwise it is called a ratio scale. An example of an interval, cardinal scale is temperature since zero Celsius and 100 Celsius are arbitrarily set based upon properties of water. The years on a calendar are another example of an interval, cardinal scale. An example of a ratio, cardinal scale is a ruler since having a value of zero implies no distance (a "true" zero). The amount of money in your bank account is another example of a ratio, cardinal scale. The cost and radiological dose attributes utilized in our value tree embody a ratio, cardinal scale.

Ordinal scales represent outcomes where the difference between outcomes does not have numerical meaning, but the ordering of outcomes implies preference to the decision maker. For example, ranking outcomes as small, medium, and large classified the outcomes, but the difference between a medium and small outcome is not quantifiable. An example of an ordinal scale would be home address numbering, where the first house on the block may be 1, the second 2, the third 3, etc., but the third house is not necessarily three times as far down the block as the first house. Surveys of preference are frequently based upon ordinal scales (I prefer apples to oranges and oranges to grapes, for example). The industrial accidents, core damage, and external attention attributes utilized in our value tree embody an ordinal scale.

Lastly, nominal scales simply list outcomes and provides the least amount of information to the decision-maker. An example of a nominal scale would be the list of fonts available from within your word processing software (Arial, Letter Gothic, Times Roman, etc.). We did not employ nominal scales within the decision analysis framework.

As part of our initial workshops, we elicited disutility information from our decision makers via application of AHP. Then, in order to determine numerical values for the attribute disutility functions, we must transform the AHP results into disutility values for each attribute scale interval. Hughes proposed a transformation that has been utilized to convert AHP values to utility (Hughes, 1986).

$$u(\text{ i'th interval }) = (v_i - \text{"worst"})/(\text{"best"} - \text{"worst"}) \qquad (9)$$

where v_i is the AHP value of the i'th interval, "worst" is the AHP value of the least desirable outcome, and "best" is the AHP value of the most desirable outcome.

Although not specifically identified in his paper, Hughes' transformation only applies to "beneficial" outcomes (measured by utility), which suggests that the transformation may not be used for disutilities unless a modification is made. Other researchers have used the Hughes transformation for disutilities (Weil and Apostolakis, 2001), but we have found that one must be careful when applying the transformation for disutilities. Consequently, we have developed a linear AHP-to-utility transformation that works for both utility (with an increasing preference function) and disutility (with a decreasing preference function). Using an example, we will demonstrate the modified transformation for disutilities.

First, assume that we operate a business where the total net worth of the business is one million euro. Now, we are asked to determine our disutility for monetary losses of three categories, zero to one million euro, one to two million euro, and two to three million euro. Since losses larger than one million euro bankrupts the company, it is presumed that the decision-maker will be indifferent between losses larger than one million euro. Also, it is highly desirable to avoid large losses in general since the decision-maker does not want the business to go bankrupt. Consequently, the response to the AHP pair-wise comparisons would look something like:

	0 to 1 M€	1 to 2 M€	2 to 3 M€
0 to 1 M€	**1**	**9**	**9**
1 to 2 M€	**1/9**	**1**	**1**
2 to 3 M€	**1/9**	**1**	**1**

where a "1" indicates indifference and a "9" indicates absolutely more important. For example, losing zero to 1 million euro is absolutely more important than losing one to two million euro.

The values from the AHP process are found (from the principal eigenvector) to be *v*(0 to 1 M euro) = 0.818, *v*(1 to 2 M euro) = 0.0911, and *v*(1 to 2 M euro) = 0.0911. Both the unaltered and the modified-Hughes transformations on the AHP values are shown via the curves in Figure 16. As can be seen in the figure, the unaltered transformation for disutility is clearly incorrect since it indicates that losses of one to two million euro have zero utility. A monetary loss of this magnitude will bankrupt the business and, therefore, should have a utility of one, not zero.

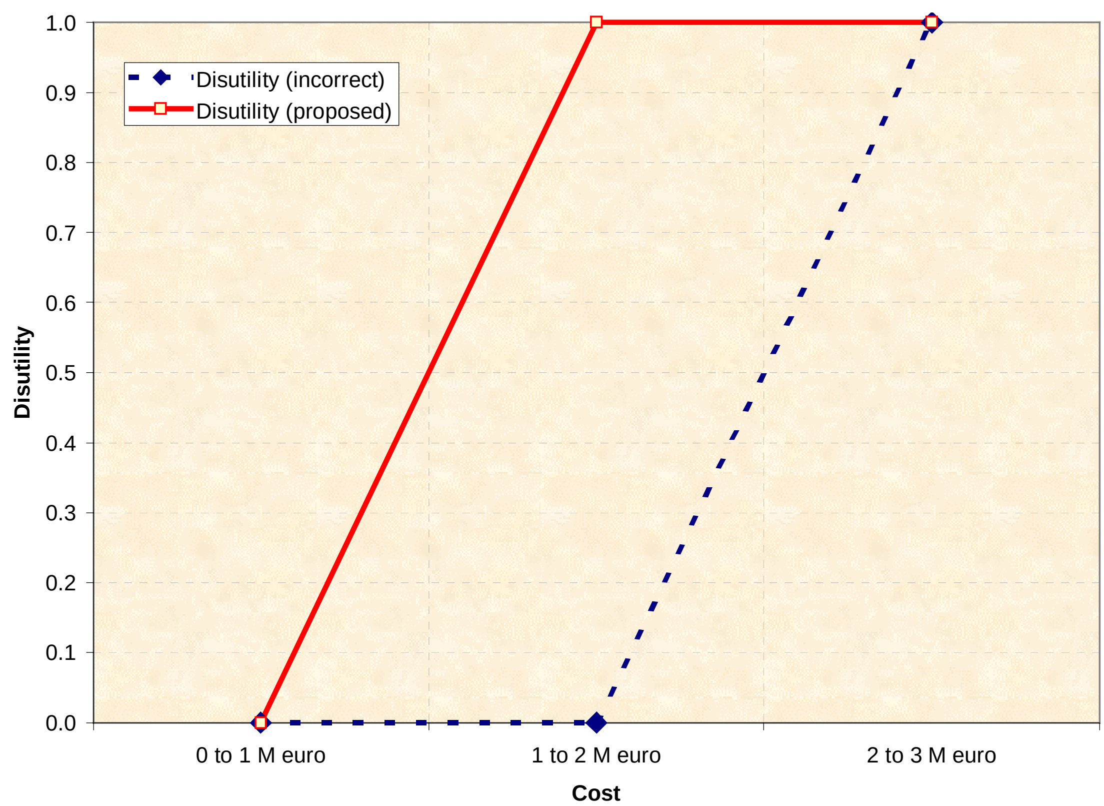

Figure 16. AHP weight to disutility transformation example.

Note that the disutility function above is plotted on an ordinal scale. If the range embodied within each scale interval is the same (which it is here, with a value of one million euro), then the shape of the curve implies risk preference. But, if the scale intervals contain non-constant values, then the shape of the curve does not necessarily indicate risk preference.

Our proposed AHP-to-utility transformation is defined by:

$$u_n = 1$$

$$u_i = u_{i+1} - \frac{v_i - v_{i+1}}{v_{high} - v_{low}} \qquad (0 < i < n) \tag{10}$$

where n indicates the number of disutility (or utility) scale intervals, i is the i'th disutility scale interval, v_i is the AHP value for the i'th disutility scale interval, v_{high} is the AHP largest value, v_{low} is the AHP smallest value, and u_i is either the disutility or utility. The number of scale intervals (i) ranges from 1 to n. Note that we could have kept the Hughes' transformation for utility and then, as a special case for disutilities, take one minus the utility in order to obtain the corresponding disutility. Instead though, the proposed transformation works equally well for either the utility or disutility since it focuses on the n'th scale as having utility or disutility of one and works back from the n'th scale, looking at incremental changes in the AHP values.

Returning to our incident management value tree, we must transform the AHP values to disutilities for all five performance measures. The initial result of this process for the cost measure is shown in Figure 17. Within this figure, we have plotted the disutility for cost on a linear, cardinal scale. Consequently, the shape of the curve will indicate the risk preference of the decision maker specific to economic losses. Note though that only four scale intervals were used to cover a large range of costs (almost ten orders of magnitude) and, consequently, one must be careful in extracting insights from this disutility function. Further, we have discovered limitations in the AHP process that restrict its use for disutility (and utility) determination. These limitations are discussed later in Section 3.

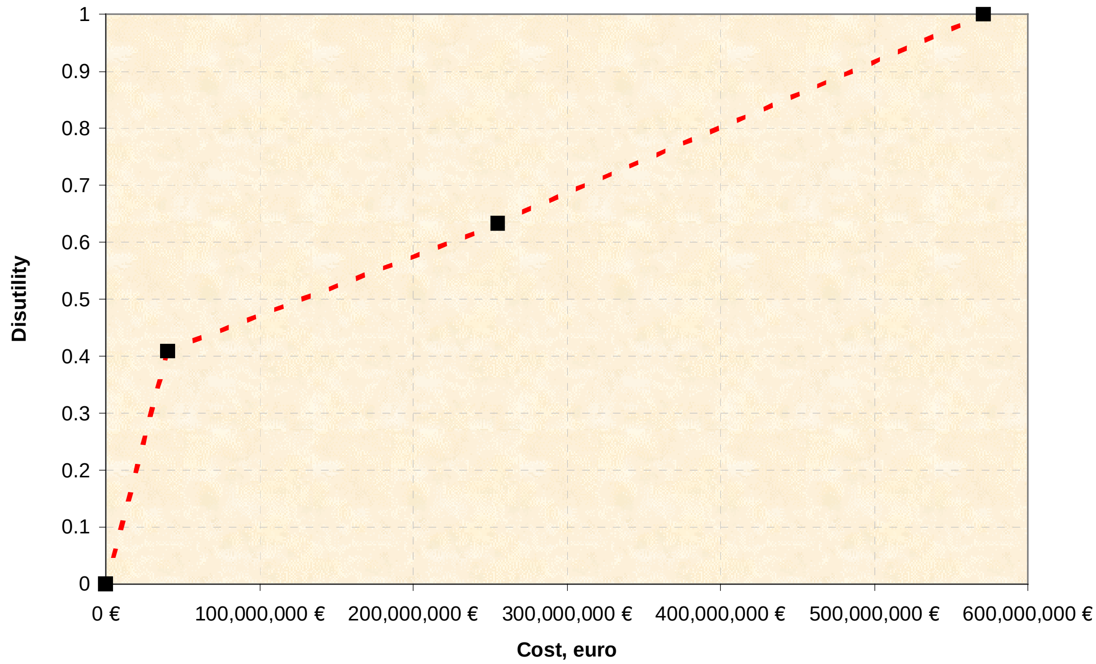

Figure 17. Example of initial determination of cost disutility function.

3.3.5 Decision Maker Risk Behavior

While the insights into the disutility function are useful, it is desirable to have more formal metrics. First, we need to determine the type of preference function we are evaluating, either increasing or decreasing in consequence. In our case, all five performance measures are ordered from best to worst outcome, so they are decreasing preference functions. Second, we need to realize the type of function, either utility or disutility. As already discussed, we are dealing with disutilities. Consequently, we are located in the lower right quadrant of the curves shown in Figure 18. From this figure, we can determine the decision maker attitude toward risk by knowing the concavity of the disutility – if the curve is concave down the decision maker is risk prone and if the curve is concave up the decision maker is risk averse.

Alternatively, the decision maker's risk behavior can be determined via a numerical evaluation of the disutility. Specifically, we would calculate (Keeney and Raiffa, 1993):

$$RB(x) = \left. \frac{\partial^2 u(x_i)}{\partial x_i^{\,2}} \middle/ \frac{\partial u(x_i)}{\partial x_i} \right. \tag{11}$$

where $RB(x)$ is the risk behavior as a function of the disutility derivatives.

From the RB(x) definition, we find:

If $RB(x) > 0$ then the decision maker is "risk prone"
If $RB(x) < 0$ then the decision maker is "risk adverse"

Also, one could use a "50-50 lottery" for the attribute outcomes to determine the decision makers risk behavior. For this query, the decision maker would be presented with two lotteries (see Figure 19). The first lottery, "A," consists of a 0.5 chance of outcome *x* and *y* (best outcomes for attributes *X* and *Y*) and a 0.5 chance of outcome *x*' and *y*' (worst outcomes for attributes *X* and *Y*). The second lottery, "B," consists of a 0.5 chance of outcome *x* and *y'* (best outcome for attribute *X* and worst outcome for *Y*) and a 0.5 chance of outcome *x'* and *y* (worst outcome for attribute *X* and best outcome for *Y*).

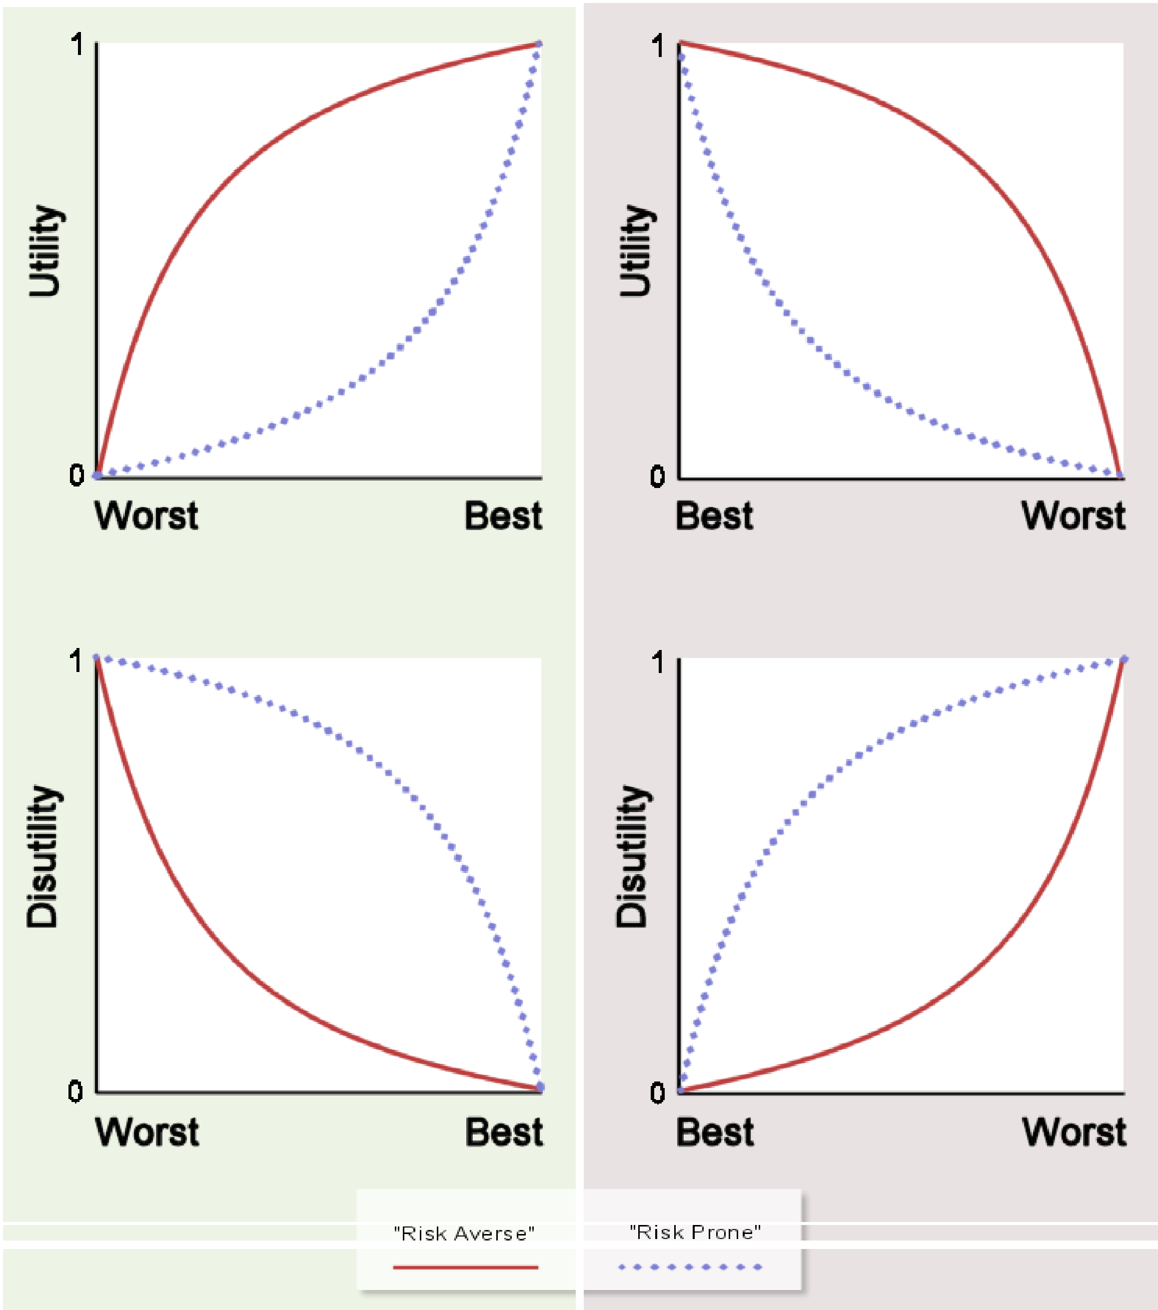

Figure 18. Illustration of “risk averse” and “risk prone” utility and disutility for increasing (worst to best) and decreasing (best to worst) preference functions.

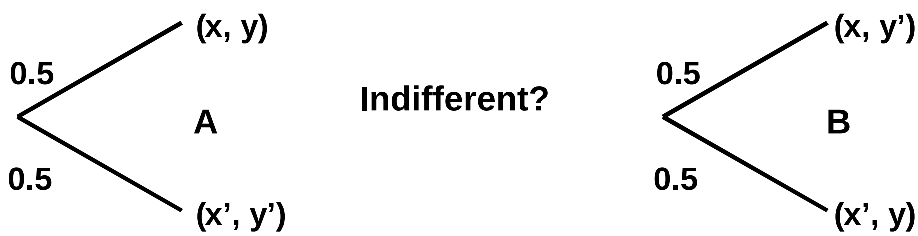

Figure 19. Lotteries used to test utility functions for decision maker risk preference.

One then determines the decision maker risk behavior by asking which lottery is preferred, or if they are of equal preference, indifference between the two lotteries. Knowledge of the outcomes is used to determine and contrast the disutility of the lottery expected value versus the disutility of the lottery itself. Based upon this information, we use the certainty equivalent criteria to determine the risk behavior:

If disutility(E[lottery]) > disutility(lottery) then the decision maker is "risk prone"
If disutility(E[lottery]) < disutility(lottery) then the decision maker is "risk averse"

Here, "risk prone" indicates decision maker prefers "gambling" on the lottery rather than taking the actual expected (i.e., average) return of the lottery. In other words, if the lottery was represented by a coin toss with a prize of $100 if heads and $0 if tails (so the expected return of the lottery is $50), then a risk prone decision maker would not accept $50 in lieu of the playing the actual lottery.

In the realm of disutility though, being risk prone is not necessarily negative. This behavior suggests that the decision maker would rather "take a chance" on getting a good outcome – while risking the possibility of bad outcome – rather than accepting a certain large (negative) outcome. Consequently, this behavior might be explained by two facts. First, the largest negative outcomes defined for our performance measures are quite extreme. For example, the cost outcome measures spans approximately nine orders of magnitude. Second, there is an inherent belief that nuclear power plants are reliable and robust – namely that the probability of an accident is quite low. Consequently, for incident management, it is unlikely that very negative outcomes will actually be realized following a decision. Thus, if the probability of an accident is low, then the decision maker may live with the risk associated with large negative outcomes. This behavior would then appear as a risk prone disutility function (Wu-Chien and Apostolakis, 1983).

From the AHP-derived curves of disutilities, we determined, via graphical methods, that our decision makers were risk prone. Also, one should note that the disutility function increases fairly rapidly over the region of interest. We were not expecting the disutility functional form to exhibit "risk prone" behavior since the general assumption towards decision makers in realistic situations are that they generally show risk adverse behavior. Nonetheless, we did find risk prone behavior for the disutilities of all five performance measures. This behavior was checked in a second workshop utilizing lottery-based questions (rather than AHP) and was found to be consistent with the AHP results (an example of the worksheet used is shown in Figure 20).

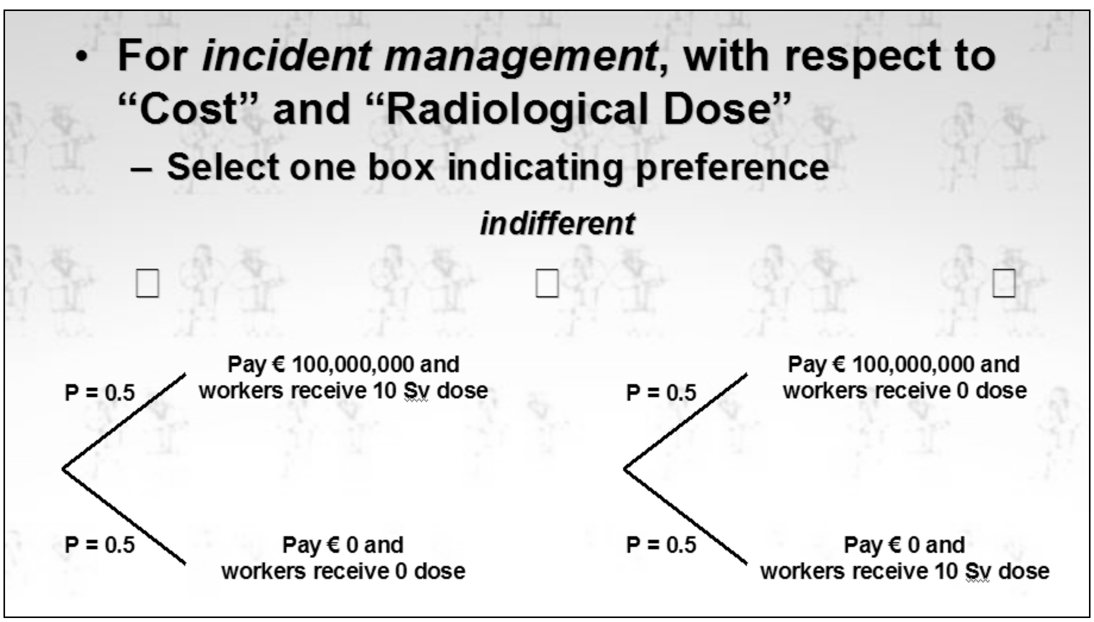

Figure 20. Example of the worksheet provided to the decision makers for determining risk preference amongst the performance measures.

If we plot the cost disutility on a logarithmic, cardinal scale, we can see the disutility function better at the low end of the outcome scale (see Figure 21). Note that in this plot, we took the geometric average of a scale interval† to get an appropriate midpoint (for example, the first interval was from zero to 14 million euro). The upper end of the disutility represents costs in the hundreds of millions of euro. This large potential outcome raises the question that small probabilities associated with the negative event may influence the decision maker preferences, even though in decision theory, the probability of an event is independent of the "worth" of that event. While competing theories such as "prospect theory" attempt to address such issues (Kahneman and Tversky, 1979), we remain committed to decision theory.

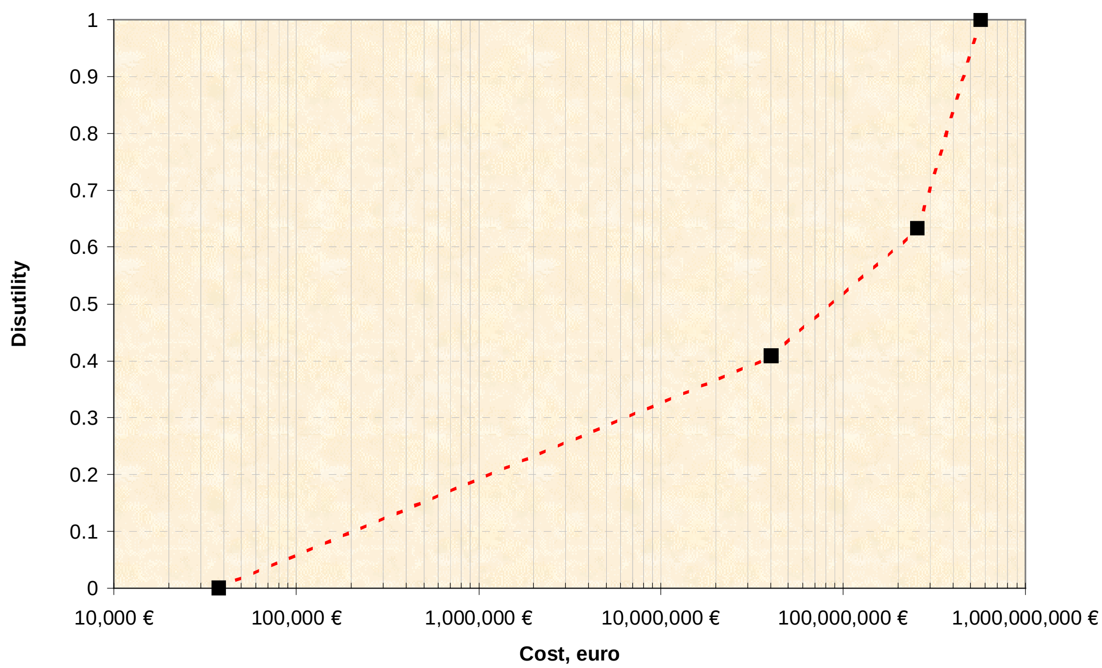

Figure 21. Example of initial cost disutility function plotted on a log scale.

---

† A geometric scale value is needed since the elicitation for both cost and radiological dose was performed by asking the importance for regions (e.g., costs of 10 to 100 million versus costs of 100 to 500 million).

3.4 Checking Consistency of the Preference Information

Upon completion of the value tree, the disutility functions, and the functional form of the performance index, the decision maker has the flexibility to perform a consistency check (also called a "sanity check") between weighted disutilities for each performance measure. Recall that the functional form of the performance index was given by a summation of the disutilities for each measure. Specifically for our case with incident management, we have:

$$\begin{aligned} PI &= w_{economics}\ u(economics) + w_{dose}\ u(dose) + w_{accidents}\ u(acc.) \\ &\quad + w_{safety}\ u(safety) + w_{stakeholders}\ u(stakeholders) \\ &= PI_{economics} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{stakeholders} \end{aligned} \tag{5}$$

Since the individual attributes are aggregated in a linear fashion, one consistency check that is available comes from the realization that a particular value of a performance index implies equivalence between any of the performance measure. For example, if we postulate a value of 0.01 for $PI_{\text{economics}}$ then, in theory, the decision maker should be indifferent to a value of 0.01 for any of the other measures. This indifference behavior arises since we have encoded preference information directly in the disutility – consequently we can utilize the additive behavior of Equation 5 to determine equivalent outcomes between the different measures.

First, let us illustrate the calculation required for the equivalence for the attributes of cost and industrial accidents. If we equate two performance measures, we are stating that the weighted disutility has the same numerical value. But, in general, the weights of any two measures will be different, thereby implying that for a given equivalence, the disutility value will change to ensure that the weighted disutilities are the same. To continue, we must then know the weights of the cost and industrial accident measures. From our decision maker, we have determined that cost has a weight of 0.32 while industrial accidents has a weight of 0.16. Now, if we choose a weighted disutility value of 0.16, this would imply that the corresponding value of the industrial accidents disutility would be maximized at 1.0 while the disutility for cost would be 0.5. Further, since we know the scales corresponding to the disutility for each measure, we have physical meaning attached to both a disutility of 1.0 on industrial accidents and 0.5 on cost. Specifically, a disutility of 1.0 on industrial accidents implies a worker fatality while a disutility of 0.5 on cost implies a loss of about 10,000,000 euro (or approximately $9,000,000 in 2002 dollars). Thus, we can state equivalence between a worker death and a loss of 10,000,000 euro based upon the value tree weights, the disutility functions, and the additive form of the weighted disutility equation. This process of picking a value of weighted disutility (0.16 in our example) is analogous to the concept of indifference curves, where for any of the five attributes with the same value of weighted disutility, the decision maker should be indifferent between the five outcomes.

We now have a method of checking the consistency of all the value tree attributes across the spectrum of ranges exhibited in each disutility function. Consequently, we evaluated all five attributes across the disutility range of 0 to 1 in order to determine applicability of the attribute equivalence. The results of this consistency check is shown in Figure 22. From this figure, one can determine the equal worth of any two attributes for a single weighted disutility value by taking a vertical slice through the attribute bars. For example, if a line were drawn from a cost of 10,000,000 euros, we would find that:

- Equivalent radiological dose is on the low end of the 50 to 2,000 mSv range
- Equivalent industrial accident is a worker fatality
- Equivalent core damage is no event
- Equivalent stakeholder interest is intervention by the regulator.

One would then want to take these equivalencies and determine if they seem adequate or not. For example, we could ask if it is reasonable that a worker fatality is “worth” about 10,000,000 euros. To answer this question, we would need to determine a societal “value of life” applicable to nuclear power plant workers and compare this monetary value to that derived from the disutility consistency check.

While the subject of assigning a monetary value to life is quite sensitive, we nonetheless found a variety of empirical evidence that provided numerical information on this subject. Not surprisingly, much of the motivation behind life valuation has come from the insurance industry. Within this industry, the main thrust behind valuation is in determining the "insurable value" rather than the value of life, where the primary difference between the two valuations is that insurable value discounts for mortality. But, the concept of life valuation does bring about a variety of definitions depending on the analysis context. Within the life insurance industry, the traditional metric of life has

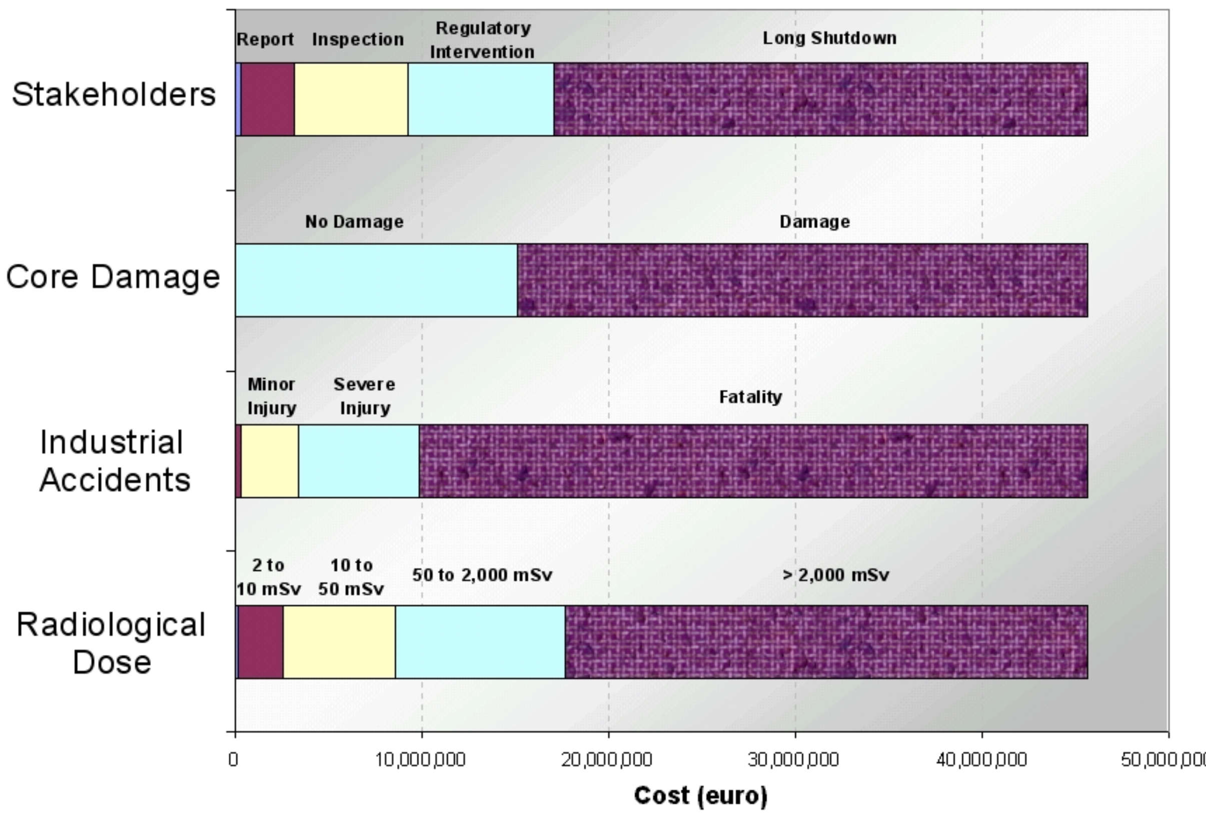

Figure 22. Consistency checks for equivalence between value tree attributes as a function of costs.

been measured by "the capitalized monetary value of future earnings" (Aponte and Dunenberg, 1968). Since the insurance industry focuses on insurable life value, it is quite concerned with specifics such as the age of the insured, health factors, earnings potential, etc. This specific approach differs from the approach required within our decision analysis framework – here we are concerned with a generalized, societal life valuation. Fortunately, this type of information is available from a variety of sources. Rather than discuss in detail the findings and results of each of these, we summarize those from the published literature, along with a brief note, in Table 4. Note that the 2002 euro amounts in Table 4 reflect standard present worth discounting (Collier and Ledbetter, 1982) using a discount rate of 3% per year. Also, at the time of this writing, the conversion from euros to dollars was approximately 1 € = $0.9.

Table 4. Summarization of select published monetary life valuation.

Value of Life (2002 euros)	Source	Notes
990,000	Hofflander, 1966	Derived from future earnings potential of "average man" in 1921.
2,600,000	Persson et al, 2001	Derived from Swedish road safety data related to a willingness to pay to improve traffic risk.
6,200,000	Keeney, 1997	Derived from mortality risks related to U.S. Government regulations.
990,000 to 8,800,000	Blomquist, 2001	Derived from "analysis of jobs with different wages and risks, consumption decisions involving changes in risk and time and money, and from direct questioning involving risk-money tradeoffs in constructed markets."
1,800,000	Burke, Aldrich, and Rasmussen, 1984	Derived from a review of "societal expenditures for life-saving safety measures."

A metric related to the question of life valuation is the NRC's cost/benefit threshold of $2000 per person-rem (U.S. NRC, 1997). If a decision alternative exhibits a value larger than $2000 per person-rem then the proposed modification is not cost effective. But, we can convert this metric to an equivalent cost-to-fatality value. Since one hundred rem equals one Sievert, and approximately seven Sievert equals a fatality, then the cost-to-fatality ratio can be calculated as $1,400,000, or approximately 1,600,000 euros. This synthesized value of life is within the data found in the general literature.

Based upon gathered evidence of the monetary worth of life, it appears that the consistency check that equates a worker fatality with a loss of 10,000,000 euros is toward the upper end of the collected data. From the data shown in Table 4, it appears that the value of life for a member of society is around two to five million euros. But, since the industrial accidents attributes focuses on worker fatality, the equivalence value should be discounted since workers typically assume a higher risk than the population at large.

We could then evaluate other consistency checks, for example between a worker fatality and a radiological dose. Performing this check leads to a dose level of approximately 100 to 200 mSv being equal to a fatality. One should contrast this check with the fact that a fatal dose (in humans) is approximately 7,000 mSv (Cember, 1992). Consequently, based upon this check, it appears that the attributes of industrial accidents and radiological dose are very inconsistent. Further checks would indicate that other attributes are similarly inconsistent. Doing this check, we see that an equivalent cost outcome, compared to a core damage event, is about 15,000,000 euros. This inconsistency yields a cost that is approximately a factor of 100 to 300 too low. For example, Rogers and Fleming (1998) estimated that a severe core damage event at an U.S. nuclear power plant could range from 5 to 12 billion euros including replacement power costs. Also, the U.S. Government mandates the maximum public liability for nuclear power plant accidents at a level of 7.3 billion dollars (6.6 billion euros) by the Price-Anderson Act (U.S. Government, 1990). Thus, a realized cost associated with a severe core damage event in the range of billions of euros is not unreasonable.

To better understand the implications associated with these types of consistency checks, we decided to evaluate other similar decision analyses that appear in the literature. While not all papers provide enough information to fully produce a consistency check, we did find a related project that did present the applicable information needed for this check. In a study of alternative siting for U.S. nuclear waste, Merkhofer and Keeney presided over a major project, funded by the U.S. Department of Energy, that evaluated proposed waste sites in a formal manner (Merkhofer and Keeney, 1987). In their study, like ours, they utilized a linear utility function to aggregate the various attributes of interest. Their functional form was slightly different though due to the method of assigning attributes weights, but the overall concept is identical as that we described earlier. In the Merkhofer and Keeney study, each attribute had a scaling constant (K) and a utility (C). Consequently, we could select an appropriate weighted utility (really disutility, but we will keep the same notation as in the paper) and compare two attributes. For example, we selected a value of 100, which implies that for a particular attribute, $K \cdot C = 100$. Now, the scaling constant for worker fatalities, public fatality (non-radiological), public fatalities (transportation related), and cost were 1, 4, 4, and 1, respectively. These weights then dictate specific values of the utility function for each attribute. For example, looking at worker fatalities, $K \cdot C = 1 \cdot C = 100$, or $C(x) = 100$. An outcome, x, of utility 100 was noted (in the Merkhofer and Keeney paper) to be equal to 30 worker deaths. Similarly, for the other attributes, we find:

- Public fatality (non-radiological) utility implies 1 death
- Public fatality (transportation) utility implies 5 deaths
- Cost utility implies $29,600,000 (in 2002 dollars, assuming a 3% discount rate)

Since a linear utility function was used in the Merkhofer and Keeney study, the outcome of each of these attributes should be equivalent to one another. In other words, 30 worker deaths should be monetarily equal to $29.6 million, or approximately $1 million per worker fatality. But, we could also equate 30 worker fatalities to one public fatality, a dubious equality indeed. Also, it is troublesome that the study allows one public fatality (non-radiological) to be equal to five public fatalities (transportation). We believe that public fatalities should be approximately equal (one-to-one) regardless of the exact nature of the fatality. Further, one can equate a single public fatality with a cost of almost $30 million, which is significantly larger than other societal public fatality costs found in the literature. Based upon the consistency problems we experienced and those from much larger projects such at that discussed by Merkhofer and Keeney, we conclude that ensuring consistency within the decision process is desirable but not easy to achieve.

3.5 Revising Preference Information to Ensure Consistency

Although we are not alone in having constructed a formal decision making framework that initially exhibited inconsistencies, we were able to address the issues related causes of the inconsistencies. If one returns to the central decision-making rule as described in Equation 5, it is evident that any inconsistencies must be resolved via modifications to either the equation itself or the parameters of the equation. Specifically, we note that three possible modifications can be made:

1. Modify the performance measure weights (as determined by the value tree)
2. Modify the functional form of the performance index equation (e.g., move to a non-linear form as described in Equation 3)
3. Modify the disutility functions themselves (e.g., rescale lower or upper endpoints of the disutilities, change the shape of the disutilities)

Of course we are free to explore modifications that consist of two or more of the items above, but as we will discuss, only the last item (3) show promise in relieving the inconsistencies that were found in our initial modeling of the decision makers preference.

3.5.1 Ensuring Consistency via Value Tree Weight Modifications

First, let us tackle the potential case of modifying the value tree performance measure weights. The existing measure weights ($w_{economics} = 0.32$, $w_{dose} = 0.16$, $w_{accidents} = 0.16$, $w_{safety} = 0.21$, and $w_{stakeholders} = 0.15$) were derived from application of AHP followed by deliberation amongst the decision makers. As such, the weights are representative of those found following the use of AHP, namely they have the characteristic that they are within a factor of ten from one another. This "bunching" of the weights from AHP is a byproduct of the rating scales, which nominally range from 1 to 9. If we construct a hypothetical AHP example to determine the extreme range of the weights (say for just two attributes) we would utilize an AHP matrix (**A**) such as:

$$\begin{array}{c|cc|} & \text{Attribute 1} & \text{Attribute 2} \\ \hline \text{Attribute 1} & 1 & 9 \\ \text{Attribute 2} & 1/9 & 1 \end{array}$$

Here, in the AHP matrix, we prefer Attribute 1 absolutely over Attribute 2. To find the weights, we need to solve the matrix equation via: $\det(\mathbf{A} - \lambda\, \mathbf{I}) = 0$, where **A** is the AHP matrix above and **I** is the identity matrix. From this equation, we can find that the normalized attribute weights are 0.9 and 0.1 for Attribute 1 and Attribute 2, respectively. Salo and Hämäläinen (1997) extended this example to the general case of an $n \times n$ matrix (i.e., n attributes) where M is the AHP upper scale numerical value (9 in our case). From this work, they noted that the maximum and minimum weights from AHP are:

$$w_{\max} = \frac{M}{n+M-1}$$
$$w_{\min} = \frac{1}{M(n-1)+1} \tag{12}$$

where M is the upper scale value (typically 9) and n is the number of attributes. We utilize these two relationships on the AHP weights to determine the permissible ratio between the maximum and minimum weights – this relationship is shown in Figure 23 where we see for the traditional AHP scale of 1 to 9, with six attributes (n = 6), the ratio of the largest weight to smallest weight is about 30. Increasing the scale from 9 to a value of 50 or 100 does increase the potential spread between the largest and smallest weight, but the principal of AHP, Saaty, strongly recommend only using the 1 to 9 scale primarily on the basis of homogeneity between the attributes (Saaty, 1980; Saaty, 1997). One of the fundamental assumptions behind AHP is "one can only compare things that are within an order of magnitude from each other" (Saaty 1997).

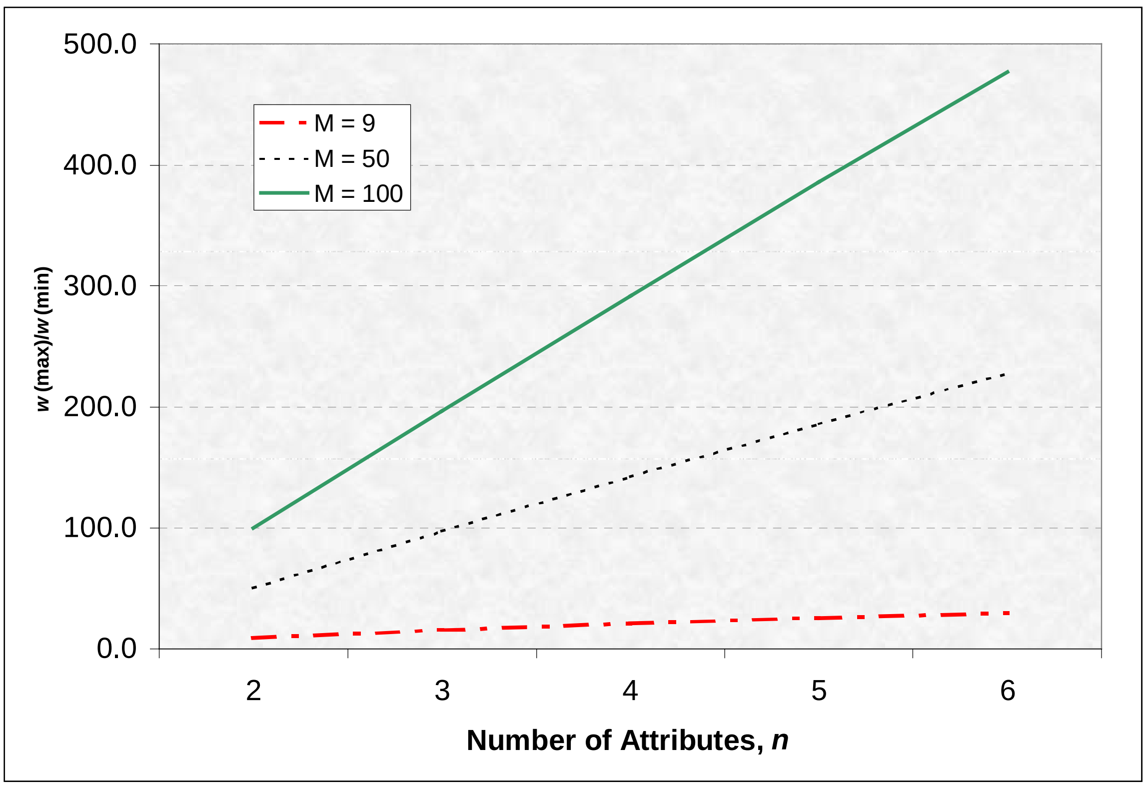

Figure 23. Illustration of the largest ratio between the maximum and minimum AHP weights as a function of both the number of overall attributes and the numerical ranking scale used for the AHP.

The example using two attributes implies that when using AHP (with the 1 to 9 scale) that the resultant weights are going to be within an order of magnitude of each other. So, what would one do if an attribute were deemed to be 100 or 1,000 times more important that another attribute? It would not be useful, in this case, to rely on AHP to assign the weights. Instead, another preference elicitation method would be required. But, in this case, one would seriously have to consider dropping the less-important (with respect to preference) attribute. A check of the preferential independence of the two attributes may indicate that dropping the lesser of the two attributes is a valid modeling choice.

Let us return then to our consistency problem. We had indicated that some of the consistency checks were off by an order of magnitude or more. Consequently, to rectify these situations, we would need to adjust the performance measure weight (specific to the consistency in question) by an order of magnitude or more. But, since the measure weights must sum to a value of one, drastically changing one measure weight may have unintended affects on other weights, thereby challenging the consistency checks of other performance measures. Also, we saw that some of the consistency checks underestimated the expected outcome, even though the attribute's utility function was set to 1.0 for the specific check. Consequently, given the current structure, maximizing the cost disutility and allowing the flexibility to change the cost performance measure weight to any value will *never* yield a reasonable cost related to a core damage event. In theory, the performance measure weights are intended to indicate preference on the attributes – they are not intended to provide a mechanism for "tuning" the results to a particular value for the performance index. So, it turns out that attempting to adjust performance measure weights is not a useful path forward towards rectifying the consistency check issues.

3.5.2 Ensuring Consistency via Non-Linear Weighted Disutility

Now, we turn to the second possible modification, namely that of going to a non-linear functional form of the performance index equation. With this particular option, it is apparent, up front, that going to a non-linear form of the performance index would greatly complicate the decision framework construction and resulting analysis. And, it is not at all evident that a more complicated form of the decision-rule equation will solve the consistency check issues. For example, again if one were to have a form of the equation that maximized the disutility for cost anytime a core damage event outcome was realized,[†] we would still only have a cost in the millions of euros rather than the (potential) hundreds of millions that could be expected. In other words, the primary problem we face with the sanity checks is in a lack of consistency between the performance measure disutility scales (e.g., the uppermost outcomes). Couple this with the fact that very few real-world applications of formal decision theory utilize non-linear forms of the decision-rule equation and we dissuaded from applying the second proposed modification. As noted by Clemen (1996) "…in extremely complicated situations with many attributes, the additive model may be a useful rough-cut approximation. It may turn out that considering the interactions among attributes is not critical to the decision at hand."

While it is generally the case that non-linear weighted utility formulations are avoided in practice, we would like to point out that, as part of the decision analysis, we do consider interactions among performance measures. But, this interaction that we will consider consists of attribute coupling via probabilistic dependence, which in turn impacts the expected value (an uncertainty) calculation. In other words, the decision rule equation given by Equation 5 is, in the end, much more complicated that it appears on the surface. For example, the PRA is used to evaluate decision options where the possibility of an accident is factored into the analysis – and, upon the realization of an accident, the probabilities and outcomes of high costs, high negative stakeholder attention, high radiological dose, and high core damage potential are all included in the overall disutility of the decision option. Thus, in effect we have non-linear characteristics related to the overall decision option ranking, but this characteristic enters into the picture via probabilities, not preferences. Note that this is similar, but not the same, as the treatment that Keeney and Raiffa (1993) give for their argument of "substitutes" related to attribute

[†] For example, the joint cost-core damage disutility would have a value of one given an outcome of a core damage event.

preference. In their discussion, they point out that the joint utility of two attributes has a weight that is a function of both attributes preference. If preferred outcomes of the two attributes lead to an even higher overall utility, then the joint attribute weight is strongly positive. But, it is possible for the joint attribute weight to be negative (since they are discussing non-linear utility functions), which would tend to lower the overall utility. And, if one attribute has a high outcome while the other has a low outcome, then the lowering of the overall utility is lessened. Consequently, if the joint attribute weight is negative, Keeney and Raiffa designate the two attributes as "substitutes" since if either one of the attributes is at a high level, then the net effect is largely positive. An example of this effect could be described by the attitude of Boston, Massachusetts, sports fans. If any *one* of the local major sports teams does well, the overall general feeling is positive…in other words, it is not necessary that *both* the New England Patriots win the Super Bowl and the Boston Red Sox win the World Series.[†]

3.5.3 Ensuring Consistency via "Measurable Equivalence"

Lastly, let us consider the third possible modification, that of modifying the disutility functions themselves. Here we are primarily concerned with making adjustments to the disutility scales, for example, by modifying the low value (best outcome) and the high value (worst outcome). We have already shown (in the two previous possible modifications) that adjusting the performance measure weights or changing the overall performance index equation is not enough to remove some of the inconsistencies we found via our indifference checks. The rational behind these shortcomings is that there is a fundamental "imbalance" in the disutility scales amongst the performance measures. Since the weights for the performance measures are within an order of magnitude from each other, this implies that both the lower scales for each measure and the upper scales for each measure must be within one order of magnitude with regard to the outcome. Let us explain this issue via a hypothetical example.

[†] The Patriots and the Red Sox are the local (to Boston) football and baseball teams, respectively.

Assume that a decision is to be made between two options, (1) being paid for a task in apples (fruit) or (2) being paid in euros (cash). Since it is desirable to use utility theory to assist in the decision process, we need a weight for apples versus euros and the utility for both apples and euros. Further assume that based upon an earlier decision process, the utility of euros was already known, where u(0 €) = 0, u(5,000 €) = 0.5, and u(10,000 €) = 1.0. Now, for the task performed, it is known that the payment in apples would be approximately one bushel,[‡] so we proceed to determine the utility of apples. Following this determination, we find:

$$\begin{aligned} &u(\text{no apples}) = 0 \\ &u(1/2 \text{ bushel}) = 0.5 \\ &u(1 \text{ bushel}) = 1.0 \end{aligned} \tag{13}$$

Also, let us assume that we prefer cash over fruit, so $w_{cash} = 0.9$ and $w_{fruit} = 0.1$. Prior to making the payment decision, we perform a consistency check. We select a weighted utility value of 0.1, which then implies a certain level of indifference between apples and euros. Specifically, we note for the two attributes:

$$\begin{aligned} &(\textit{cash}) \quad 0.1 = w_{cash} \cdot u(x_{cash}) = 0.9 \cdot u(x_{cash}) \\ &(\textit{fruit}) \quad 0.1 = w_{fruit} \cdot u(x_{fruit}) = 0.1 \cdot u(x_{fruit}) \end{aligned} \tag{14}$$

By knowing the utility functions for both euros and apples, we can determine the equivalencies between these two attributes. From the utilities defined in the previous paragraph, we would find that x_{cash} = 1,110 euros and x_{fruit} = 1 bushel of apples. But, it is obvious that one bushel of apples should not be equivalent to 1,110 euros. Instead, visiting the grocery store will show that one bushel of apples can be purchased (at retail prices) for approximately 50 euros. Thus, our consistency check indicates that the indifference between euros and apples is off by about a factor 22. We could then ask,

‡ One bushel of apples contains approximately 100 apples and weights about 19 kilograms.

"what is the driver behind this discrepancy?" The answer here, as is the same for our incident management case, is that the utility scales are not consistent. If the maximum outcome for the fruit attribute is one bushel of apples, the maximum outcome for the cash attribute should be no more than an order of magnitude larger than the cash value of a bushel of apples. Otherwise, with inconsistent utility scales, we experience equalities between attribute utility that simply do not fit reality.

In the end, to fix the inequalities that we saw with our initial disutilities, we utilized a process that we call "measurable equivalence." Typical methods for disutility determination include not only AHP but also "certainty equivalent" and "lottery equivalent" (Clemen, 1996). The certainty equivalent approach asks the decision maker to compare a lottery of outcomes against a known, certain quantity. But, due to biases in comparing uncertain outcomes against certain outcomes, researchers have proposed the lottery method whereby the decision maker compares two different lotteries. We propose a fourth method, that of measurable equivalence, whereby facts are used to adjust disutilities. For example, in the apple example above, a measurable equivalence that can be used to determine the utility is the fact that apples cost about 50 euro per bushel. We can utilize this known relationship to adjust either the cost or apple disutility function such that the consistency checks provide reasonable results. This method is similar to the "portable genius" approach described by Howard (1971) in that the domain constraints should, to some extent, focus the decision problem. Ultimately, we utilized the method of measurable equivalence to finalize the disutility functions for the decision advisor prototype.

3.6 Limitations in the AHP Methodology

Due to the difficulties we experienced in using AHP for the disutility determination, we decided to collect additional preference information from our decision makers. In a follow-on workshop to the original AHP exercises, we elicited from our decision makers preference information on the five performance measures. But, in this second case, we decomposed the scales into two intervals for some of the performance measures in order to obtain preference details for select regions. For example, we chose to decompose the cost measure into two sub-intervals rather than attempting to span nine orders of magnitude in a single AHP elicitation. This decomposition does increase the information collection burden, but it does provide additional data points across the measure scale.

In order to bring the two sub-intervals into a single disutility function (in the case of the cost performance measure), we "joined" the regions together at their end-points (or "pivot" points). Specifically, the upper end-point of the first region was equated to the lower end-point of the second region. This process of joining sub-intervals via pivots utilized the normalized AHP values (that is, prior to the transformation to the disutility values) which are obtained from the principal AHP eigenvector (Saaty and Vargas, 2000). Consequently, each disutility value in the first interval (from 0 to 750,000 euro) was normalized by the ratio of the 750,000 euro AHP value of the second interval over the 750,000 euro AHP value from the first interval. An example of the AHP sub-interval step for the cost performance measure is shown in Table 5. The AHP value normalization factor that was used is 0.59/0.048, or 12.3.

Table 5. AHP and disutility data for the cost performance measure by joining two sub-interval at their end-points.

Interval	Cost (euro)	AHP Value (raw)	AHP Value (normalized)	Disutility
1	1,000	0.45	5.53	0.00
	5,000	0.36	4.40	0.21
	50,000	0.15	1.83	0.67
	750,000	0.048	0.59	0.90
2	*750,000*	0.59	*0.59*	0.90
	10,000,000	0.27	*0.27*	0.96
	100,000,000	0.11	*0.11*	0.97
	1,000,000,000	0.032	*0.032*	1.00

From the disutility values in Table 5, we can construct a disutility function for the cost performance measure. We show the result of this preliminary function in Figure 24. Like the previous AHP-derived curves, this cost disutility exhibits a risk prone attitude. But, one should notice that the cost disutility rises rapidly. For example, the cost curve shows a disutility of 0.5 around a cost of 20,000 euro. The question then is this a reasonable value for a disutility of 0.5? One method of answering such a question is through the use of the "certainty equivalent."

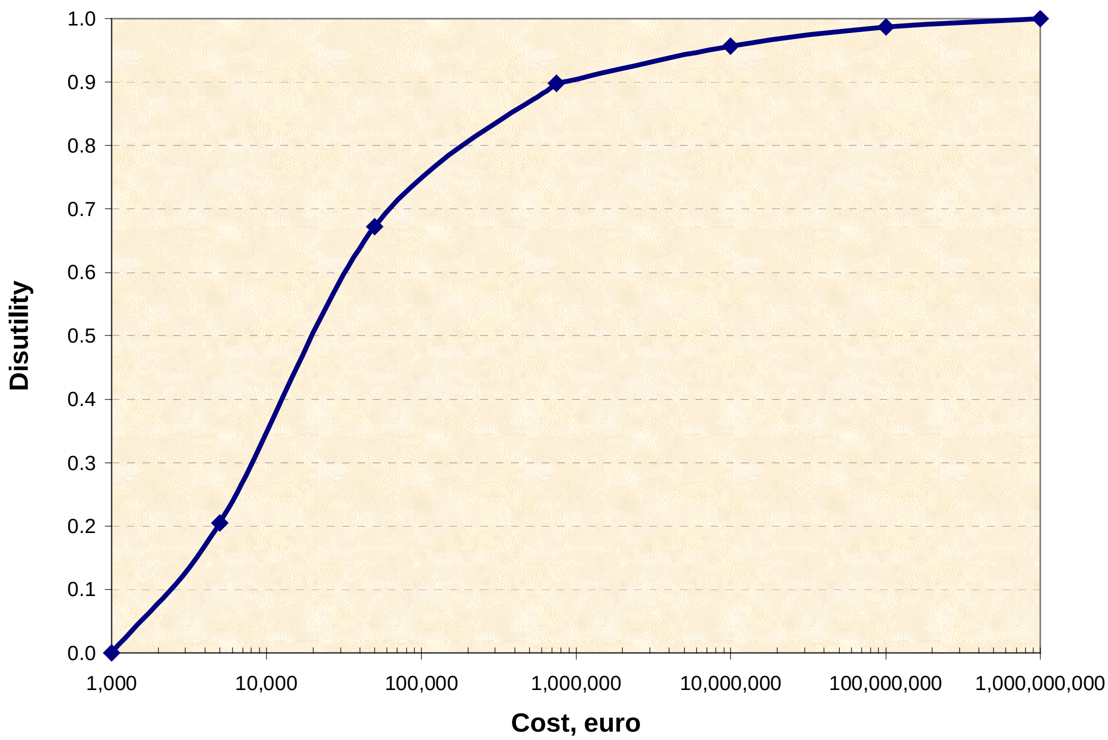

Figure 24. Pivot-based AHP disutility curve for the cost performance measure.

The definition of the certainty equivalent is tied to the concepts of lotteries. For a performance measure *X*, the expected outcome of a lottery of *n* number of *x* outcomes, each with a unique probability *p*, is given by

$$E[x] = \sum_{i=1}^{n} p_i x_i \tag{15}$$

The expected utility of the lottery of performance measure *X* is then

$$E[u(x)] = \sum_{i=1}^{n} p_i u(x_i) \tag{16}$$

For a lottery based upon a coin toss, a so-called 50-50 lottery, the expected value of the lottery is calculated as E[*x*] = (0.5 x_1) + (0.5 x_2), where x_1 and x_2 are the two outcomes of interest, respectively. The certainty equivalent for a lottery is the outcome of the performance measure *X*, denoted by *x'*, where the decision maker becomes indifferent between *x'* for certain or the lottery itself. Since preference is expressed via utilities, the certainty equivalent may be expressed as a function of utilities (Keeney and Raiffa, 1993)

$$x' = u^{-1}\left(\sum_{i=1}^{n} p_i u(x_i)\right) \tag{17}$$

Note that the certainty equivalent value, *x'*, is expressed in terms of the units for the *X* measure. For example, if *X* is cost, then *x'* is found in terms of monetary units. For any measure though, the certainty equivalent may be determined once the utility is known for that performance measure. Further note that the literature frequently describes the process of determining an equivalence in terms of a 50-50 lottery, but the concept is valid for any combination of *n* probabilities, where the only constraint is that $\Sigma p_i = 1$. Of course the value x' changes as the lottery probabilities are adjusted, simply due to the fact that the expected utility of the lottery changes (as noted in Equation 16).

To illustrate the certainty equivalent calculation, we will utilize the pivot-based cost performance measure disutility. We will use a 50-50 lottery for two cases, first for the case of outcomes zero and 1 billion euro and second for the case of outcomes zero and 20,000 euro. In the first case, 1 billion euro has a disutility of 1.0, so the certainty equivalent value is found by

$$x' = u^{-1}\left(0.5u(0) + 0.5u(1{,}000{,}000{,}000\,euro)\right) = u^{-1}(0.5) = 20{,}000\,euro \tag{18}$$

The second case is calculated via a similar equation. The certainty equivalent result for both cases are shown in Figure 25, where the certainty equivalent for the first case is approximately 20,000 euro while, in the second case, the certainty equivalent is approximately 6,000 euro. These two values are noted by the vertical lines on the plot.

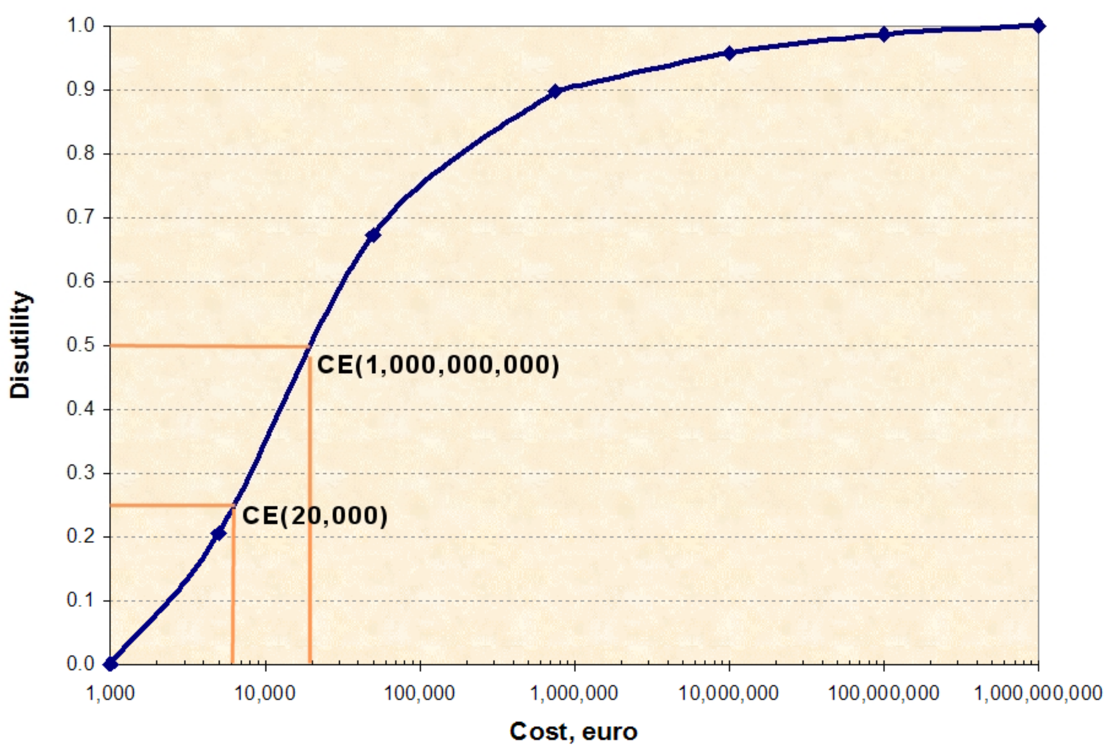

Figure 25. Certainty equivalent calculation for the pivot-based cost disutility.

We propose to use this concept of certainty equivalent as a check on the disutility function itself. Using the intervals defined in Table 5, we plot, in Figure 26, the certainty equivalent (of a 50-50 lottery) for interval 1 (costs below 750,000 euro), interval 2 (costs above 750,000 euro), and the joined intervals (via pivots). We also show the certainty equivalent, as a function of cost, in the case of a "risk neutral" decision maker. In this figure, if the certainty equivalent is less than the risk neutral curve at that same cost, then the associated disutility exhibits risk prone behavior. From the relationships described in Figure 26, we point out an interesting observation – namely that the pivot approach captures the certainty equivalent behavior *only* for the lowest interval. The *pivot approach does a poor job* of representing the certainty equivalent for larger cost outcomes. For example, the pivot approach indicates that the certainty equivalent for a 50-50 lottery (where the maximum loss is one billion euro) is only 20,000 euro (as we also illustrated in Figure 25).

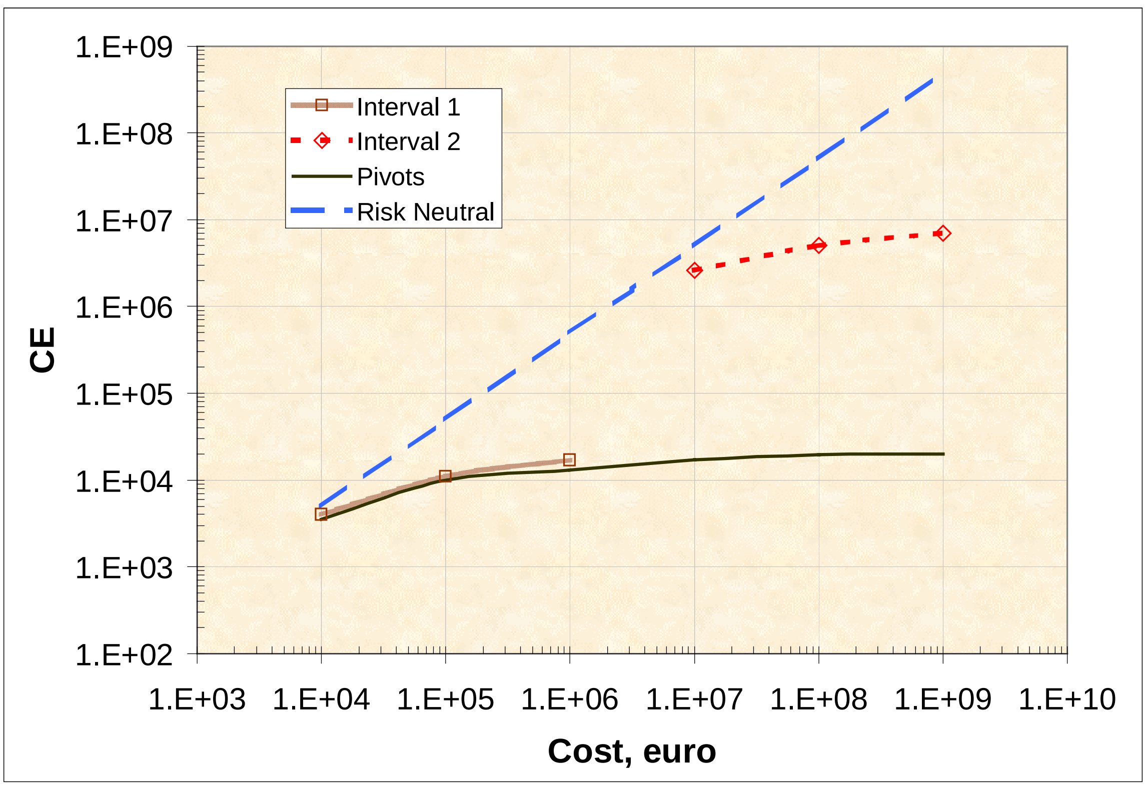

Figure 26. Certainty equivalent as a function of cost from the AHP elicitation.

A certainty equivalent of 20,000 euro on a one billion euro, 50-50 lottery is far too small and indicates that the decision maker is far too risk prone, or to be more precise, that the disutility as indicated by AHP is much to risk prone. The fact that the pivot approach suggested by Saaty (Saaty and Vargas, 2000) does not preserve the certainty equivalent embodied in the performance measure upper intervals indicates it probably should not be utilized for the elicitation of disutility when the potential outcomes span a wide range.

The AHP pivot approach for disutility is not valid if one desires to maintain the certainty equivalent toward the upper end of the performance measure. The undesired features of the AHP pivot method include:

1. The pivot-based disutility implies that the decision maker is very risk prone. The disutility is not near the "risk neutral" curve for much of the performance measure range and may differ by orders of magnitude from the risk neutral curve.

2. The pivot-based disutility preserves the certainty equivalent only for the low cost regions (e.g., less than one million euro). In higher regions, the certainty equivalent is suspect. For example, the AHP results suggest that the decision maker would only pay 20,000 euro to avoid a "50-50" lottery where the two outcomes are either lose nothing or lose 1 billion euro. In reality, the decision makers would pay much more than 20,000 euro to avoid large losses of this type.

From the determination of the certainty equivalent for the pivot case, it became apparent that the disutility via AHP was driven, to a large extent, by the initial cost intervals. Consequently, we became concerned that even for a "non-pivot" AHP application, that the resulting disutility function may be driven by information specific to the initial portion performance measure. To check this, we then evaluated cost preference using AHP, where the span was from 100,000 to 100 million euro. The data that we had collected from our decision makers for this interval is shown in Table 6.

Table 6. AHP and disutility data for the cost performance for the interval of 100,000 to 100,000,000 euro.

Cost (euro)	AHP Matrix				AHP Value (raw)	AHP Value (normalized)	Disutility
100,000	1	***4***	6	8	0.876	0.567	0.00
1,000,000	1/***4***	1	6	8	0.448	0.290	0.52
10,000,000	1/6	1/6	1	8	0.171	0.111	0.85
100,000,000	1/8	1/8	1/8	1	0.051	0.033	1.00

To test our hypothesis that AHP results are driven by the initial scale region, we adjusted the AHP matrix concerning preference on losses of 100,000 euro versus losses of 1 million euro.[†] We modified the "4" value in region I of the AHP matrix (Table 6), first to a value of "1," where we called this case "no impact" to indicate that the decision maker is indifferent between costs of 100,000 and 1,000,000 euro. Second, we modified the same value in region I of the matrix to a value of "2," where we called this case "minor impact." The results of these AHP calculation are shown in Figure 27, where we plot the corresponding disutility function for the cases described above. In addition to the AHP-derived cases, we plot the disutility that would be realized for a risk neutral decision maker.

As we see in Figure 27, the disutility curve changes significantly depending on the value in region I of the matrix. This degree of sensitivity indicates that the AHP process is driven by the initial regions specified for the problem domain. Since the number and magnitude of these initial regions are somewhat arbitrary (for example, instead of having a single interval from 100,000 to 1,000,000 euro, we could have five regions within this

[†] Examples of these losses at a nuclear power plant is a shutdown of approximately one day or less.

span), it appears that the AHP-derived disutility is somewhat arbitrary. In other words, it is possible to obtain dramatically different functional forms for disutility depending on the nuances of setting up the AHP scales. Unfortunately, this feature could be used to drive the disutility functional forms to almost any desired outcome. For example, if one evaluates the "no impact" curve shown in Figure 27, one would notice that the initial portion of this curve demonstrates "risk averse" behavior (below one million euro) while upper portions of the curve are "risk prone." So, the decision maker preference (as indicated by the AHP curve) may vary depending on the AHP implementation, which should not be the case in practice. The decision maker preference on a performance measure outcome should be independent (in theory) of the measure scale quantities used to derive the disutility. Consequently, it appears that using AHP to determine disutility violates Utility Axiom 2 (equivalence).

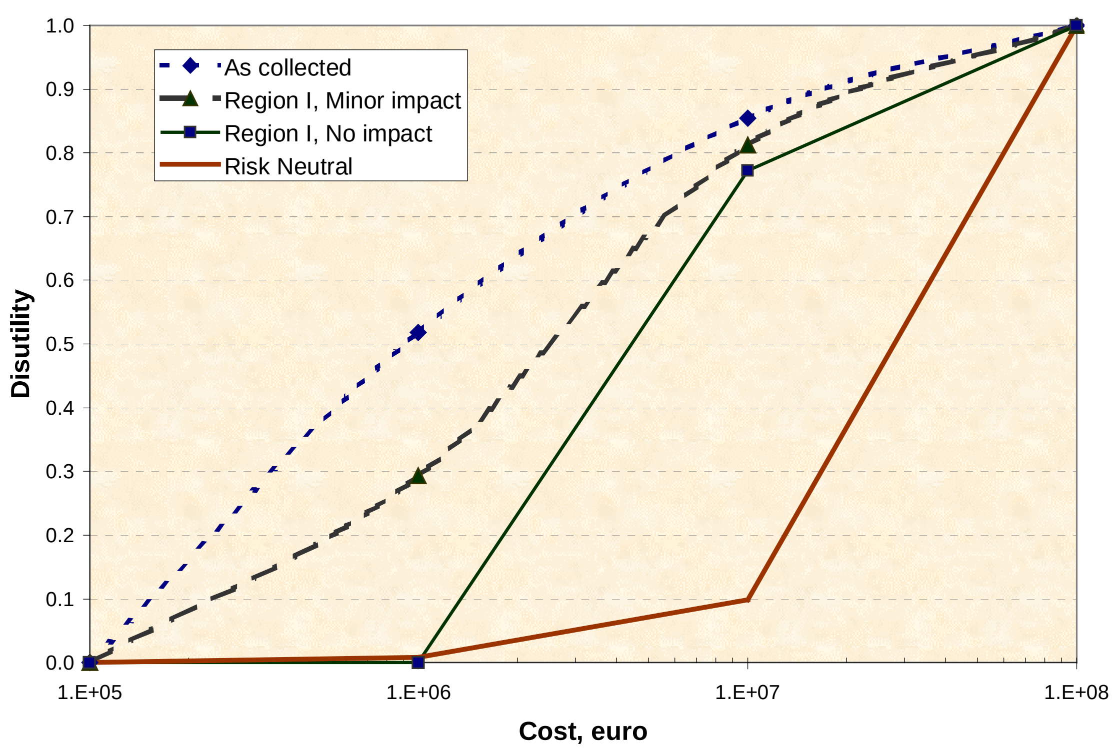

Figure 27. Sensitivity of the disutility curve for the cost performance to changes in one region of the AHP matrix.

As a final investigation to the AHP-to-disutility determination, we decided to see if there was any correlation between the performance measure scale range (as measured by the orders of magnitude between the low and upper end-points) and the inconsistency metric of the AHP matrix. For this analysis, we utilized the "consistency index" formulation proposed by Saaty (1980), which is defined as:

$$CI = \frac{\lambda_{\max} - n}{n-1} \tag{19}$$

where $\lambda_{\max}$ is the largest eigenvalue from the AHP and n is the number of performance measure scale points.

We looked at the potential for AHP inconsistency correlation for both the cost and radiological dose measures. The results of the cost analysis is shown in Figure 28 while that for the radiological dose is shown in Figure 29. We can clearly see that as the range increases (i.e., as the performance measure scale spans larger values), the AHP inconsistency index increases. Generally, CI values above 0.10 to 0.15 are considered to be "inconsistent" and would provide motivation for reconsideration of the AHP matrix. This correlation between the AHP scale and the inconsistency is further evidence for the need to utilize an alternative method (other than AHP) for disutilities.

While our original motivation behind using AHP for disutilities was its ease of use and increasing acceptance in the realm of decision making, we also desire to have a method that is robust, believable, and passes our "sanity" checks over the range of consequences of concern. For our application though, where disutility spans several orders of magnitude, we have become critical of using AHP as a disutility-generating mechanism. Consequently, for the disutilities that are to be used in the decision advisor prototype, we chose not to use AHP.

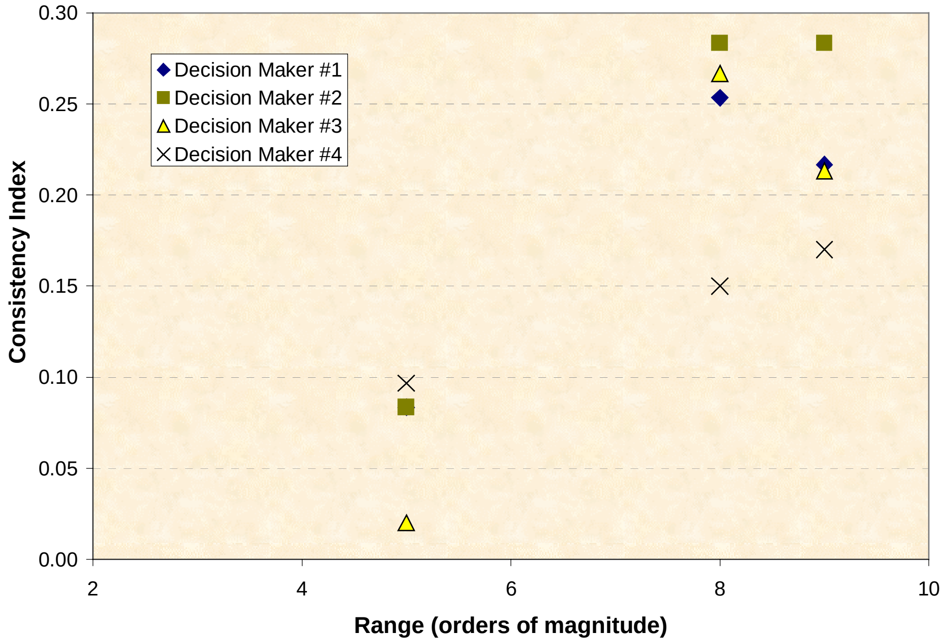

Figure 28. The AHP consistency index (for four decision makers) as a function of the cost magnitude (as measured by the scale orders of magnitude).

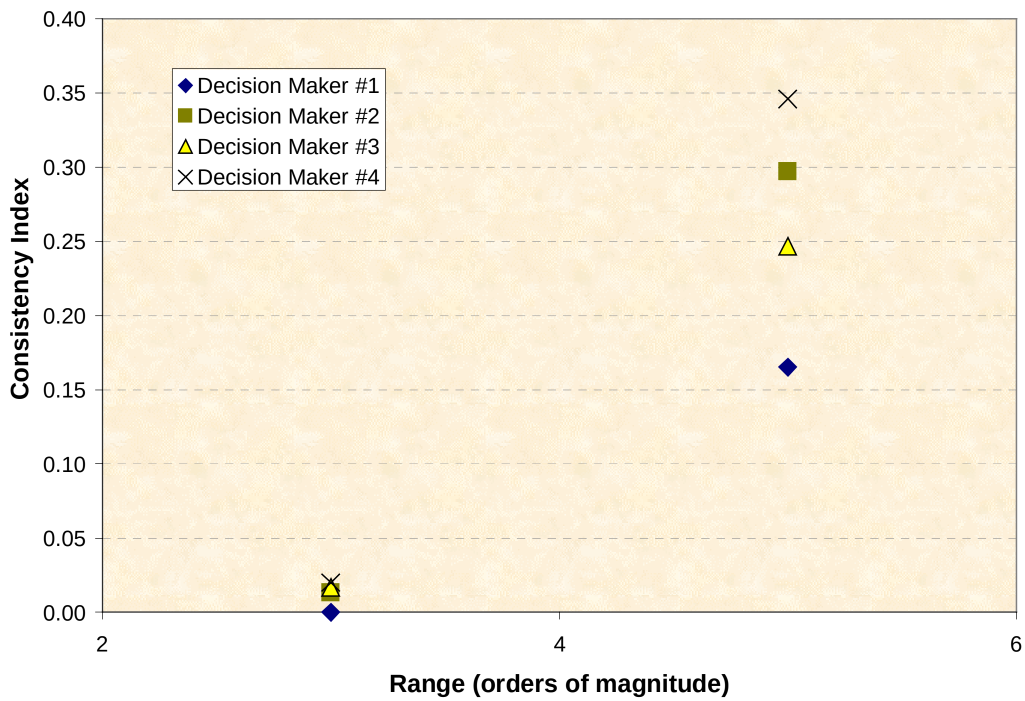

Figure 29. The AHP consistency index (for four decision makers) as a function of the radiological dose magnitude (as measured by scale orders of magnitude).

3.7 Finalizing the Disutility Functions

In order to modify the disutility scales for the incident management framework, we must consider two guiding principles of the adjustment. First, we need to ensure that the maximum outcome (the worst case) for each performance measure has, within reason, about the same level of "consequence." This first principle is employed to ensure that like outcomes are used up-front in the preference elicitation process. Second, where possible, the performance measure indifference points should be constrained by actual measurable equalities. Here, we are defining a "measurable equality" as two performance measure outcomes that, if put side-by-side in reality, would be equivalent. Thus, this second principle is used to bring real data into the decision process while simultaneously helping to reduce the subjectivity present when utilizing preference information. As an example of the types of measurable equalities one might utilize in the context of incident management, we tabulated several for various performance measures in Table 7.

We should point out that not every decision analysis domain has measurable equalities. For example, the frequently used textbook problem of "buying a car" has attributes such as fuel economy, cost, and style as fundamental to the purchase of a car. But, in this case, there would not be natural equivalencies between any two of the attributes. So, in one respect, we are fortunate to have several measurable equalities for the problem of incident management, but we still need to resolve the issue of the extreme range in disutility.

Table 7. Examples of measurable equalities between the incident management performance measures.

Performance Measures	Equivalence Conditions
Cost and Industrial Accidents	The value of life cost associated with a worker fatality is approximately 2 million euros
Industrial Accidents and Radiological Dose	A fatality is realized upon exposure to a dose of approximately 7 Sv
Cost and Core Damage	An estimated cost associated with an accident resulting in partial core damage (but no vessel breach) is approximately 750 million euros
Cost and External Attention	Regulatory intervention that forces the plant to suspend operations costs 333,000 euros per 24 hour period

3.7.1 Determining the Upper Endpoint for the Disutility

At this point, we can revisit each of the attributes in order to force the disutility scales to be consistent within the value tree framework, including weights, described in Figure 15. For each disutility, the initial end point represents the best outcome and has value of zero. This outcome is defined by the lack of any negative impact for the attribute (no costs, no dose, no external attention, etc.). The other end of the range for each disutility is the worst outcome and has a value of one. This outcome is defined by the maximal situation that could occur within the context of incident management. But, it is this end point that must be made consistent between the five attributes.

In order to ensure consistency, we must first find the attribute outcome, of the five, that currently represents the worst outcome. To determine this "anchor," we show the scales that were used for the initial disutility elicitation in Table 8. Then, one must focus on the last row in the table, specifically the information pertaining to disutility equal to one. Within this row, it is plausible to assume that the worst outcome of all five “worst” outcomes is that of core damage. Consequently, we must recast the other four performance measures (cost, industrial accidents, radiological dose, and external attention) such that they are comparable (within ratios of the attribute weights). But, since we are reconstituting the disutility scales, we are also going to allow constraints based upon measurable equalities (for example, see Table 7).

As we can see in Table 8, inconsistencies are likely to exist due to the discrepancy in the performance measure scales. For example, the upper end of the industrial accidents disutility scale is a fatality, but for the cost scale we see a maximum outcome of hundreds of millions of euros. Consequently, it is evident that we need to adjust the disutility scales, either by increasing the performance measures that are too low or by decreasing the ones that are high. Once might question the need to have plant safety end in core damage if the focus is on incident management. But, this outcome was preserved due to the desire expressed by the decision maker.

Table 8. The scales used in the initial determination of performance measure disutility.

Range	Cost (millions of euros)	Industrial Accidents	Radiological Dose (Sv)	Core Damage	External Attention
Best (u = 0)	< 14.3	No injury	< 0.002	None	None
	14.3 to 114	Minor injury	0.002 to 0.01		Report event
			0.01 to 0.05		Inspection
	114 to 571	Severe injury	0.05 to 2		Regulatory intervention
Worst (u = 1)	> 571	One fatality	2	Core damage	Long shutdown

For the five performance measures that we require, we need to determine appropriate "mappings" that can translate disutility information from one measure into another. These measurable equalities, or mappings, should consist of observable or readily available data, preferably collected from actual experience. Once these measurable equalities are identified, we will use them to transform disutility for the five performance measure functions, where we must account for the difference in performance weights that were identified as part of the value tree determination (see Figure 15). This scaling process is similar to the consistency checks that were performed earlier for the AHP-derived disutilities where we will force equivalent outcomes to have equal weighted PI (Equation 5). The general steps for the disutility application of measurable equivalence are:

1. Denote appropriate measurable equivalence for the performance measures.
2. Determine which performance measure has the largest "outcome."
3. Set the upper-bound of the remaining performance measures based upon Step 2.
4. Determine, perhaps by using lottery equivalence techniques, the disutility function for one of the performance measures, preferably for the measure that has the largest number of equivalencies identified in Step 1.
5. Determine the remaining disutility functions by equating outcomes via the weighted PI, starting with the function identified in Step 4.
6. Deliberate in order to complete the disutility for regions or points not assigned by use of the previous five steps.

We now demonstrate these steps though the application for our five performance measures. First, we need to identify applicable measurable equivalencies. While we noted a few example equalities earlier in Table 7, we determined a total of ten that were of use. This final set of equalities is listed in Table 9 along with their references. We also display a plot of select equalities in Figure 30.

Table 9. Measurable equalities that were utilized to determine the final disutility functions for the five performance measures.

Performance Measure					Equality	Reference
Cost	Dose	Attention	Core Damage	Accidents		
●	●				222,000 euro = 1 Sv	U.S. NRC, 1997
	●			●	7 Sv = 1 Fatality	Cember, 1992
●				●	1 Fatality = 2E6 euro	Table 4
●				●	1 Injury = 2E5 euro	Burke, Aldrich, and Rasmussen, 1984
			●	●	Core damage < 10 fatalities	Burke, Aldrich, and Rasmussen, 1984
		●	●		Core damage = long shutdown	Burke, Aldrich, and Rasmussen, 1984
●			●		Core damage = 7.5E8 euro	Burke, Aldrich, and Rasmussen, 1984
●		●			Regulatory report = 3E5 euro	Inferred from our decision maker as a likely outcome
		●		●	Fatality = inspection	Inferred from our decision maker as a likely outcome
●		●			1 day outage = 3.3E5 euro	Estimated by our decision maker for actual plant operation losses

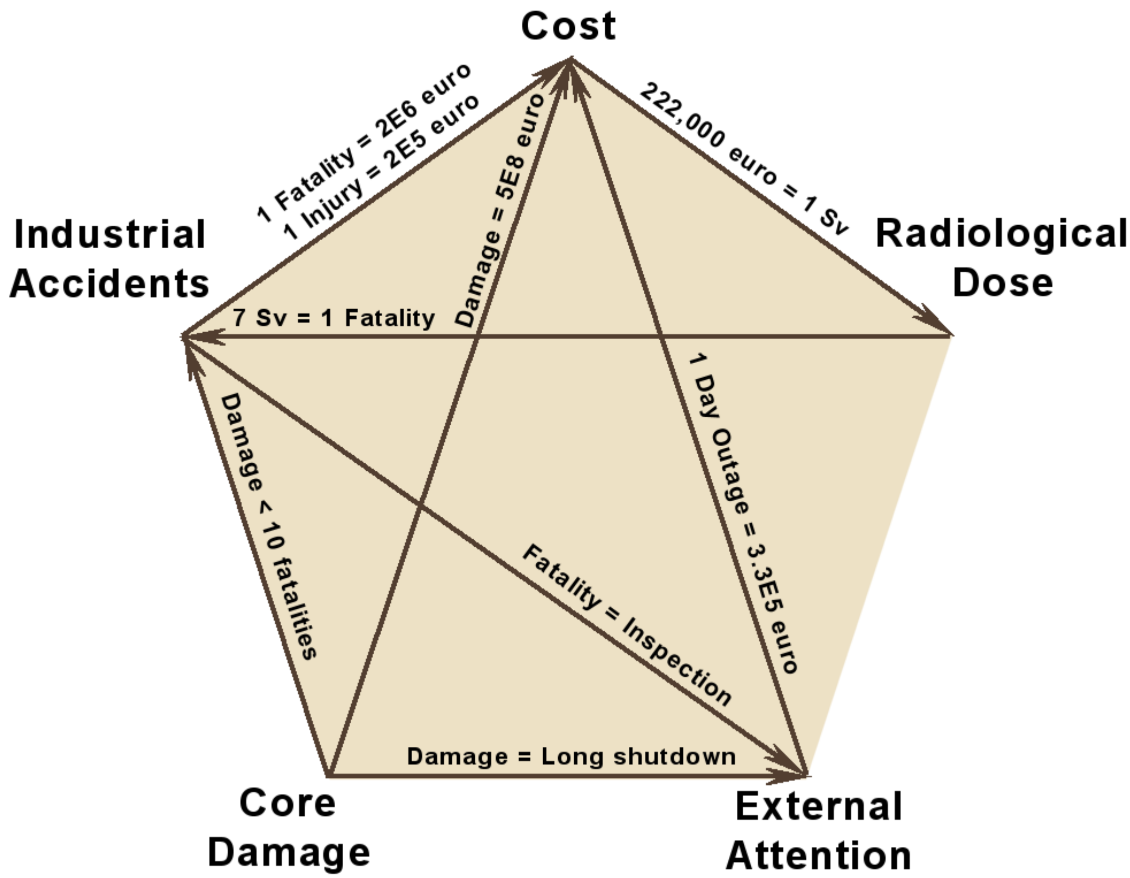

Figure 30. Measurable equivalence factors used to develop the final disutility functions. The arrows indicate the "direction" of influence from one measure to another (e.g., receiving a dose of 7 Sv induces a fatality).

Previously, we discussed the fact that the performance measure of core damage embodies our largest "outcome," namely that of a core damage event. While we *are not* focusing on severe events where radioactive material is released to the environment, minor core damage events are of interest. These types of events fall close to the "category II" accidents evaluated by Burke, Aldrich, and Rasmussen (1984), which they classify as "core damage events without pressure vessel melt-through, no significant release of radioactive material, and no offsite public health impacts or property damage."

Nonetheless, these events constitute our bounding event, both in consequence and in (low) probability.

We now need to scale the upper bound for each of the four remaining performance measure disutilities. This scaling is accomplished by utilizing: (a) an applicable measurable equivalence and (b) the ratio of the performance measure weights. Let us demonstrate the technique to determine the cost disutility upper bound.

For the cost disutility, we have an equality that indicates a core damage event is worth (approximately) 750 million euro. Also, the performance measure weight for cost is 0.32, while it is 0.21 for core damage. Thus, the equivalent core damage to *maximum cost* is found by returning to the PI equation and equating these two measures:

$$0.32\,u(Cost_{\max}) = 0.21\,u(Core\ Damage) \qquad (20)$$

Since the disutility for the maximum event is defined to be 1.0, we know that u(*Core Damage*) = 1. Even though the actual cost of a core damage event is around 750 million euro, this value can not be the maximum cost for the disutility due to the fact that the cost weight is 50% higher than the core damage weight. But, we can find the maximum cost (which has a disutility of 1.0) from the cost-to-core damage equality, or:

$$Cost_{\max} = u_{CD}^{-1}(1.0) \times \left(\frac{0.21}{0.32}\right) \times \left(\frac{Cost}{Core\,Damage}\right) \qquad (21)$$

where $u_{CD}^{-1}(1.0)$ is the inverse disutility given a maximum value of 1.0 (which is just a core damage outcome) and *Cost/Core Damage* is the actual core damage cost per event (750 million euro). This calculation yields a maximum cost of 500 million euro. Note that this transformation is performed at this point to determine the maximum scale value for cost. We do not perform a linear transformation of the entire cost function based upon the relationship described in Equation 21. Instead, we use the individual measurable equivalence relationship to determine specific points of indifference on a particular disutility curve. Then, once enough points are described, we may fill in the remaining points via deliberation or, in the case of cardinal-scale based disutilities, may fit the points to an appropriate disutility function. Note that in some ordinal-scaled based cases that each category could be determined via the measurable equivalence technique if enough relationships were defined.

In order to determine the remaining disutility points, we need to have information concerning the cost disutility function. In order to determine the cost disutility, we can not rely on AHP, so we utilized a worksheet based upon lottery equivalencies. We asked our decision makers to compare two lotteries, each with different economic losses. For example, one of the questions was posed as:

Lottery A = 0.5 * (100,000,000 euro cost) + 0.5 (no cost)

Lottery B = p * (1,000,000,000 euro) + (1 - p) * (no cost)

What value of p makes these two lotteries indifferent? p = _____?

Once the decision maker specifies their value of p, in effect they have given us cost disutility values due to the application of *Coherence Axiom 6*, that of preference. For example, in the case above, we are allowed to write:

$$0.5u(100\ million\ euro) = p \times u(1\ billion\ euro) \tag{22}$$

where we note that u(no cost) = 0. But, since u(1 billion euro) = 1, we have:

$$u(100\ million\ euro) = 2p \tag{23}$$

In this case, our decision makers specified that p was equivalent to 5.5E-3, so we now have another point on our cost disutility curve, namely u(100 million euro) = 1.1E-2. In total, we elicited four points on the disutility curve, ranging from costs of 100,000 euro up to 100 million euro. Joined with these four values are the two additional points, the lower (u = 0) and upper (u = 1), for a total of six disutility points. This set of cost disutility is shown in Table 10.

Table 10. Cost disutility points based upon lottery comparisons.

Cost Outcome (euro)	Disutility
0	0.0
100,000	4.5E-6
1,000,000	3.0E-5
10,000,000	1.7E-3
100,000,000	1.1E-2
500,000,000	1.0

We then fit these points to a few disutility function types (logarithmic, exponential, polynomial). For the cost disutility points, we found that a 3rd-order polynomial fit the points well, with only an overestimation of the disutility toward the low end of the scale. This plot is shown in Figure 31. Further deliberation by the decision makers could provided adjustments to this disutility cost curve, but we assume that it provides an adequate representation of the decision maker preferences.

Now that we have the cost disutility function, we can determine the dose, industrial accidents, and external attention upper endpoints since we also have measurable equivalence for each. The other three maximums are found like before. For dose:

$$Dose_{\max} = u_{Cost}^{-1}\left(\frac{0.16}{0.32}\right)\times\left(\frac{1\,Sv}{222{,}000\,euro}\right) \tag{24}$$

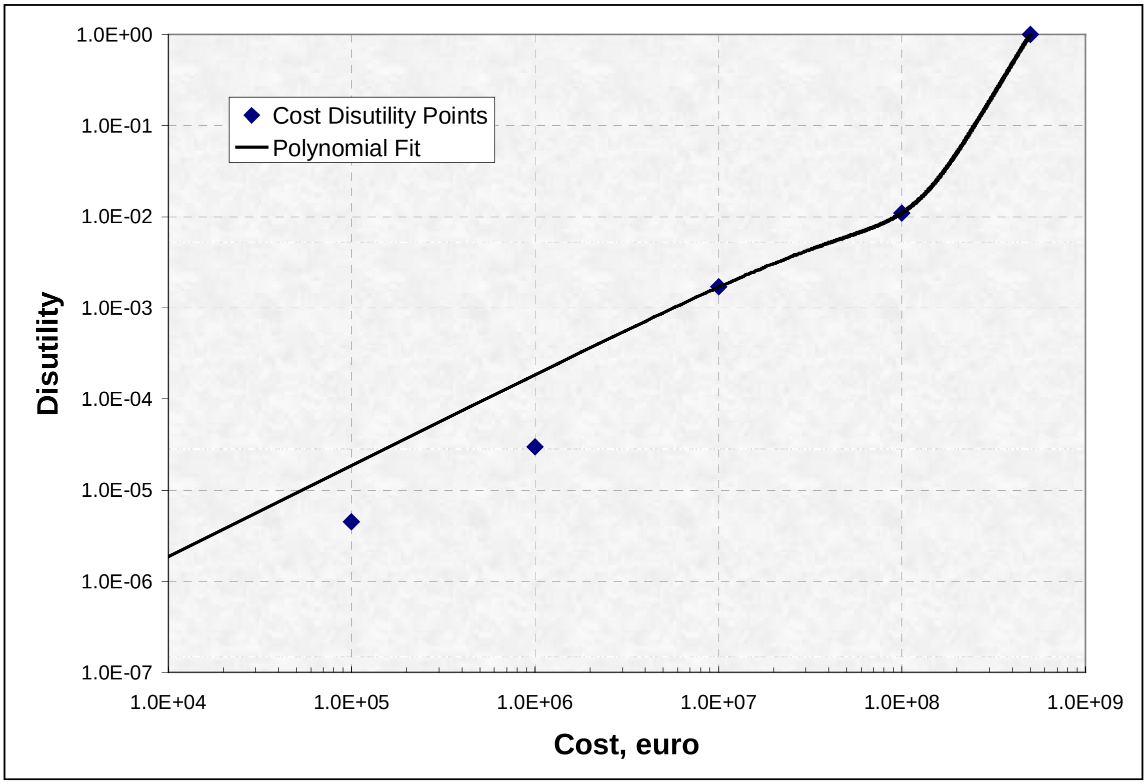

Figure 31. Final disutility curve for the *cost* performance measure that is used in the incident management prototype (log-log scale).

where $u_{Cost}^{-1}(0.16/0.32)$ is the inverse disutility for cost and 1 *Sv/*222,000 *euro* is determined by knowing the measurable equality shown in Table 9. From our function for cost, $u_{Cost}^{-1}(0.16/0.32) = u_{Cost}^{-1}(0.5)$ = 427 million euro. Thus, the maximum value on the dose disutility scale is:

$$Dose_{\max} = 427{,}000{,}000\,euro \times \left(\frac{1\,Sv}{222{,}000\,euro}\right) \approx 2000\,Sv \qquad (25)$$

Again, note that this we have now only described a single point on the disutility dose curve. We do not apply this transformation to the entire curve, but instead describe as many points as possible using the measurable equivalence technique. But first, let us determine the maximum values for the remaining disutility curves.

For the industrial accidents disutility, we find the maximum scale value as:

$$Accidents_{\max} = u_{\cos t}^{-1}(0.16/0.32) \times \left(\frac{1\ fatality}{2{,}000{,}000\, euro} \right) \approx 200\ fatalities \qquad (26)$$

While the maximum scale for accidents embodies an extremely large number of fatalities, we note that the scale for this performance measure is a ordinal-based one. Consequently, we are going to describe the largest region on the scale as any situation that involves ten or more fatalities. Thus, this region (> 10 fatalities) will have a disutility of one, which bounds the case of 200 (or more) fatalities.

For the external attention disutility, we find the maximum scale value as:

$$Attention_{\max} = u_{\cos t}^{-1}\left(\frac{0.15}{0.32} \right) \times \left(\frac{1 day\ shutdown}{333{,}000\, euro} \right) = 1260\, days \approx 3.5 years \qquad (27)$$

since u_{Cost}^{-1}(0.15/0.32) = 419 million euro. Thus, the maximum regulatory attention has an outcome of a 3.5 year shutdown. This maximum, and the maximums for all five performance measures, are listed in Table 11.

Table 11. Performance measures weights and maximum outcomes.

Performance Measure	Weight	Maximum Outcome for Disutility
Cost	0.32	500 million euro
Radiological Dose	0.16	2,000 Sv
Industrial Accidents	0.16	>10 fatalities
Core Damage	0.21	Core Damage
External Attention	0.15	3.5 year shutdown

3.7.2 Filling in the Disutility Functions via Measurable Equivalence

The remaining measurable equivalencies are used to determine additional points on the radiological dose, industrial accidents, and external attention disutilities. For example, we have an equality between a lethal dose (7 Sv) and the cost of a fatality (2,000,000 euro). Based upon this, we can determine another point directly on the dose disutility curve, since:

$$0.16 \times u(7\,Sv) = 0.32 \times u(2{,}000{,}000\,euro) \tag{28}$$

The point we are looking for is the disutility at the 7 Sv level, or:

$$u(7\,Sv) = \left(\frac{0.32}{0.16}\right) \times u(2{,}000{,}000\,euro) \tag{29}$$

But, from the cost disutility, we find *u*(2,000,000 euro) = 3.9E-4. Thus, *u*(7 Sv) = 7.8E-4. Also, we have the relationship between 1 Sv and 222,000 euro, so *u*(1 Sv) = 8.8E-5 since u(222,000 euro) = 4.4E-5. We now have four points on the dose disutility curve. Again, we fit the points (shown in Table 12) to several curves, where the closest fit was found to be a 2nd-order polynomial. The radiological dose disutility curve is shown in Figure 32.

Table 12. Dose disutility points based upon measurable equivalence.

Dose Outcome (Sv)	Disutility
0	0.0
1	8.8E-5
7	7.8E-4
2,000	1.0

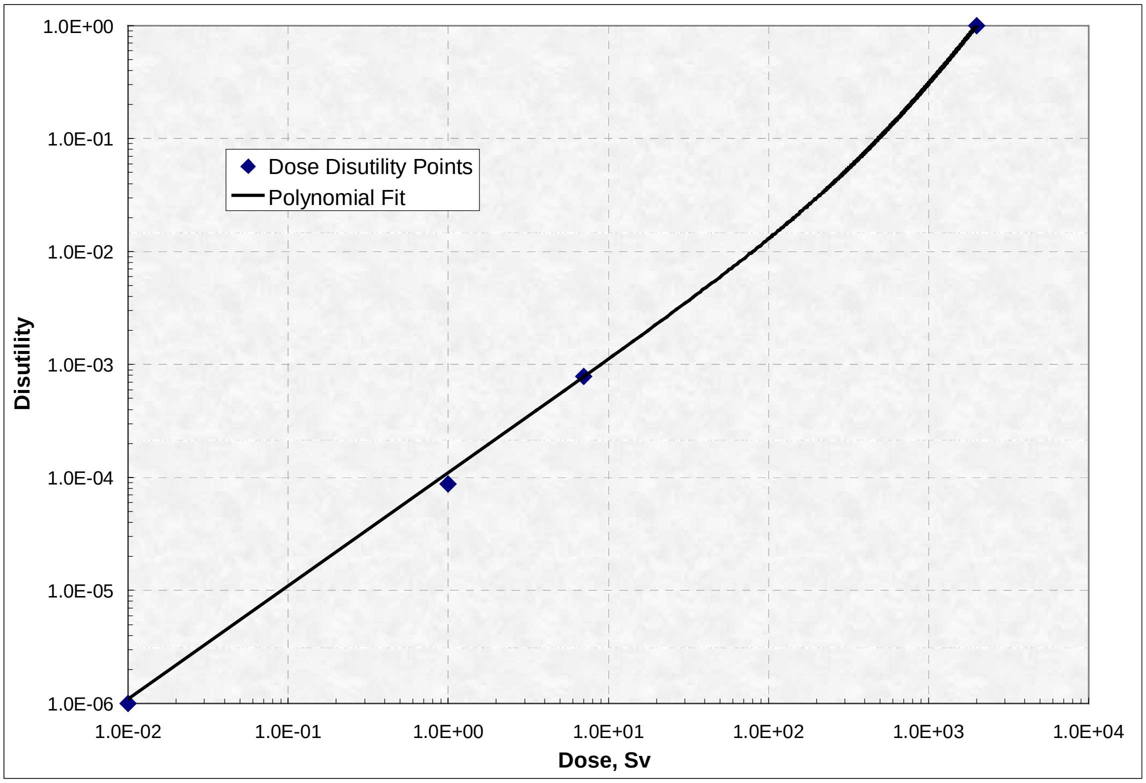

Figure 32. Final disutility curve for the *radiological dose* performance measure that is used in the incident management prototype (log-log scale).

Now, we can determine points for the industrial accident disutility curve. We know that there is equivalence between a major injury and a cost of 200,000 euro. From this, we can determine:

$$0.16 \times u(major\ injury) = 0.32 \times u(200{,}000\ euro) \tag{30}$$

Thus, we can now find the disutility for a major injury, or:

$$u(major\ injury) = \left(\frac{0.32}{0.16}\right) \times u(200{,}000\ euro) = 8 \times 10^{-5} \tag{31}$$

We have the relationship between a fatality and a 7 Sv dose, so *u*(fatality) = 7.8E-4 since u(7 Sv dose) = 7.8E-4 and the dose and accident weights are the same. As an extension to the single fatality category, we also utilize an interval of up to 10 fatalities, were we assume that a cost associated with this outcome would be approximately times the single fatality case. Thus, *u*(2 to 10 fatalities) = 6.6E-3 since u(20,000,000 euro) = 3.3E-3.

The only remaining point on the industrial accident disutility curve is that for the case of a minor injury. We did not have an identified equivalence for this outcome. Consequently, we assumed that a minor injury was approximately equivalent to 20,000 euro. Thus, *u*(minor injury) = 4.0E-6 since u(20,000 euro) = 2.0E-6.

We now have the points (shown in Table 13) required to determine the industrial accident disutility function. The industrial accident disutility function is shown in Figure 33.

Table 13. Industrial accident disutility points based upon measurable equivalence.

Accident Outcome	Disutility
None	0.0
Minor injury	4.0E-6
Major injury	8.0E-5
Single fatality	7.8E-4
2 to 10 fatalities	6.6E-3
> 10 fatalities	1.0

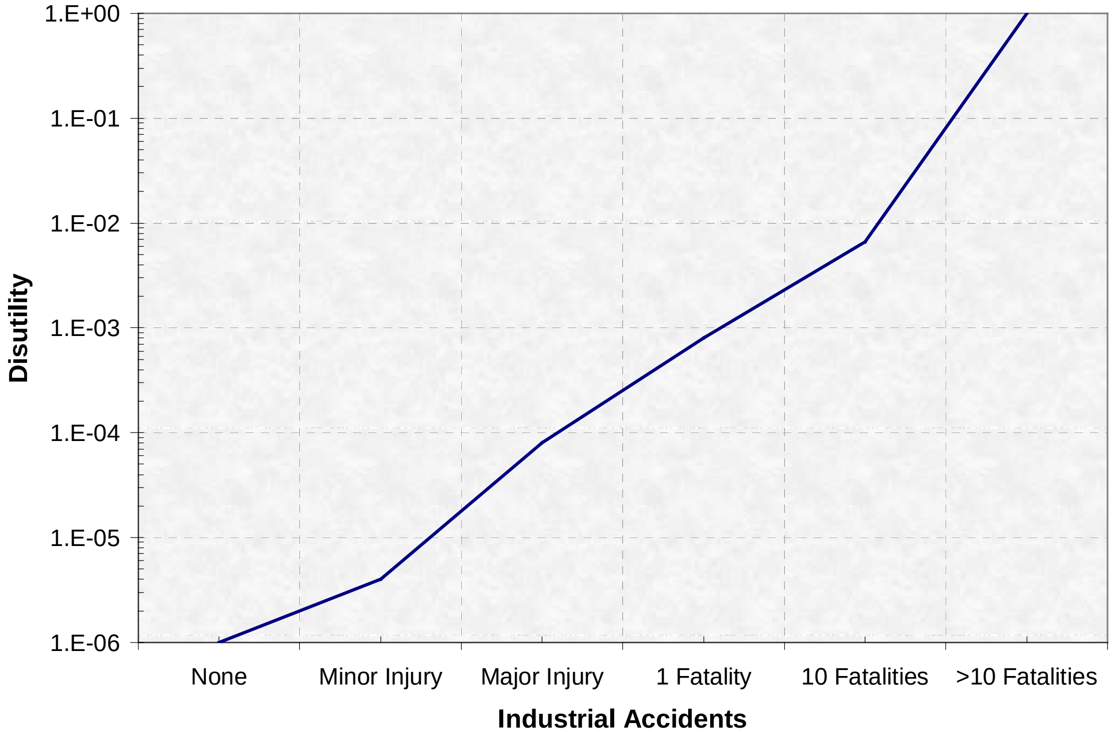

Figure 33. Final disutility curve for the *industrial accident* performance measure that is used in the incident management prototype.

Lastly, we can determine points for the external attention disutility curve. We know that there is equivalence between a report and a cost of 333,000 euro. From this, we can determine:

$$0.15 \times u(report) = 0.32 \times u(333{,}000\, euro) \tag{32}$$

Thus, we can now find the disutility for a report to the regulator, or:

$$u(report) = \left(\frac{0.32}{0.15}\right) \times u(333{,}000\, euro) = 1.4 \times 10^{-4} \tag{33}$$

Other equivalencies that we utilized for the external attention disutility points include the assumption that an inspection equates to an economic burden of 2.6 million euro. Also, we assumed that the impact from regulatory intervention was between that of an inspection and a one-year shutdown. From these assumptions, we can calculate the points (shown in Table 14) required to determine the external attention disutility function. The external attention disutility function is shown in Figure 34.

Table 14. External attention disutility points based upon measurable equivalence.

Attention Outcome	Disutility
None	0.0
Report	1.4E-4
Inspection	1.1E-3
Regulatory intervention	4.6E-3
Shutdown 1 year	1.9E-2
Shutdown 3.5 years	1.0

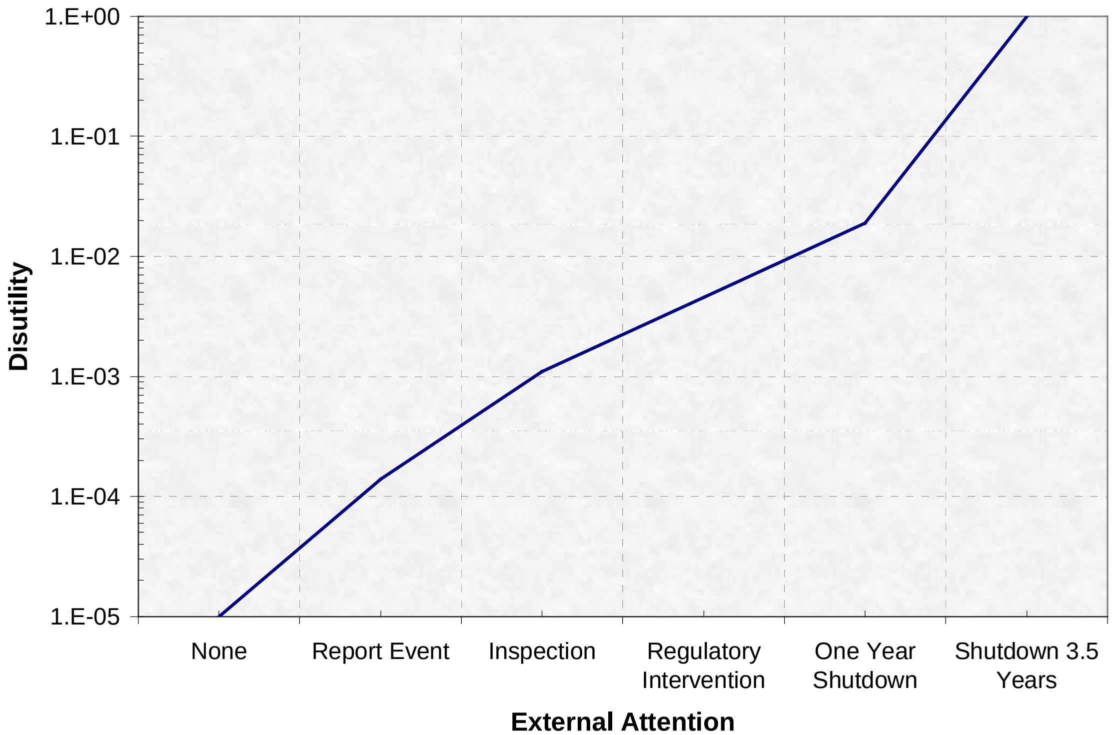

Figure 34. Final disutility curve for the *external attention* performance measure that is used in the incident management prototype.

The one remaining disutility curve is that for plant safety. But, since there are only two outcomes for this performance measure, either core damage or not, the disutility function is straightforward. This function is shown in Figure 35.

The process of measurable equivalence relies on a robust set of performance measure weights. As discussed in Section 3.3.3, we used AHP to derive the performance weights. We believe that AHP is suitable for this application since the performance measures are "within an order of magnitude" from each other and the performance measures are single entities (e.g., cost versus dose), unlike the scales that we saw for disutility. Thus, we will not have the problem of identifying intervals within the performance measure. Instead, the AHP application for performance measure weights is a one-to-one comparison of measures (within the context of incident management).

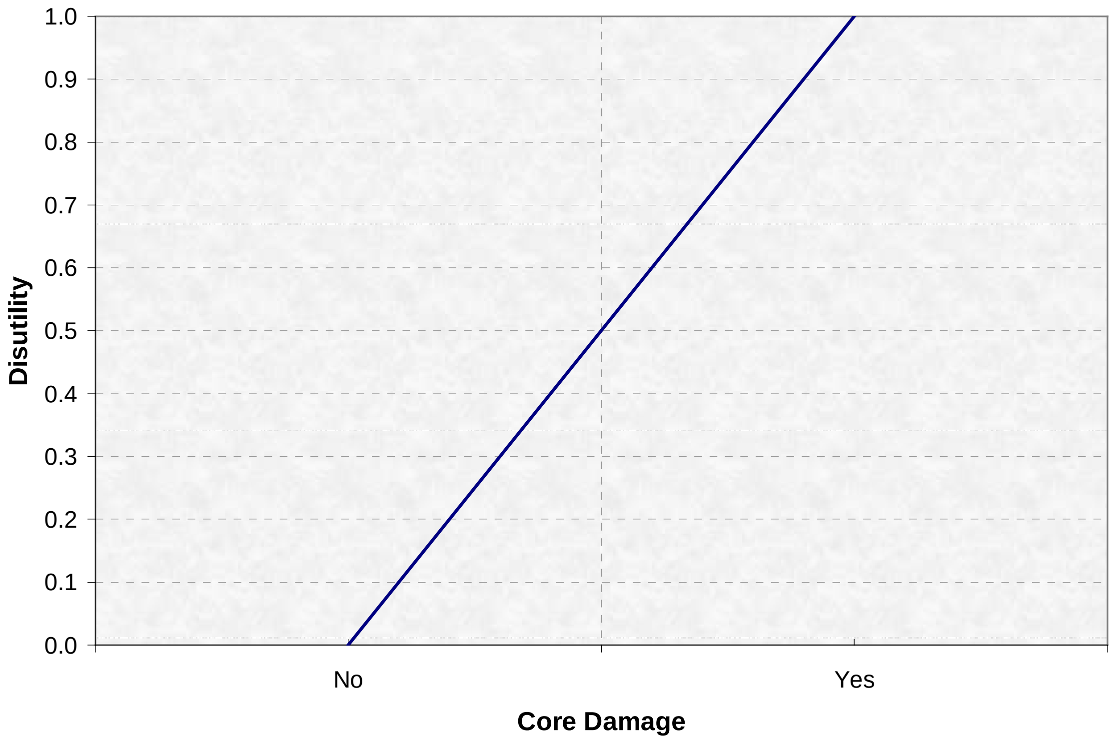

Figure 35. Final disutility curve for the *plant safety* performance measure that is used in the incident management prototype.

3.7.3 Consistency Checks of the Disutility Functions

Once the disutilities were determined via the measurable equivalence technique, we still need to validate the curves by performing consistency checks and determining the certainty equivalent. First, let us look at the consistency (or sanity) checks between the performance measures. We plot the weighted PI comparisons for each of the performance measures in Figure 36. Evaluating this figure, we see that the consistency between performance measures is quite good. For example, we see that a fatality is equal to approximately 2 million euro, which is also equal to a dose of about 7 Sv. We should not be surprised at this consistency since equalities between performance measures are used to determine that disutilities. A list of equality comparisons is also provided in Table 15 which contrasts the disutility results obtained by way of AHP (see Figure 22) and that by measurable equivalence. In every case, the measurable equivalence results are superior to those obtained via AHP.

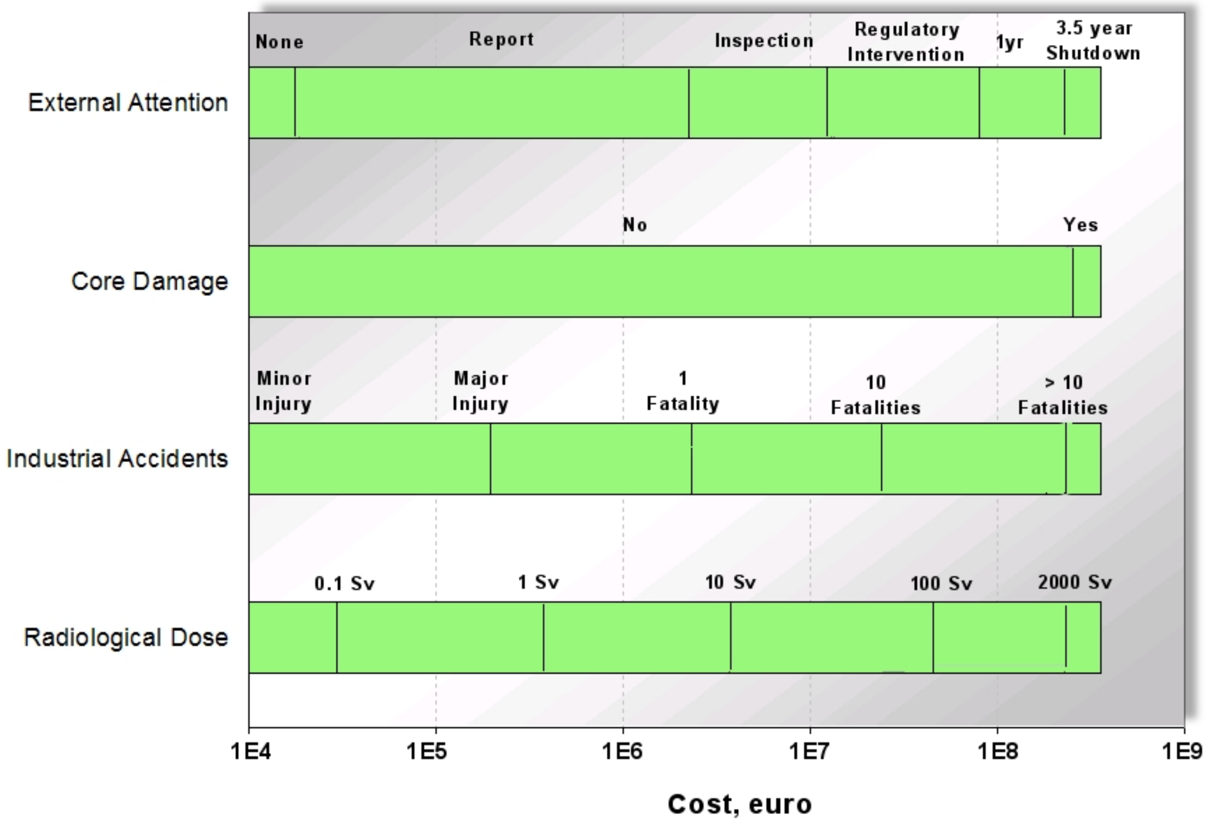

Figure 36. Consistency checks, plotted as a function of cost, for the performance measures when using the measurable equivalence method.

Table 15. Example consistency check comparison between AHP-derived and measurable equivalence-derived disutilities.

Consistency Check	Original AHP Results	Measurable Equivalence Results
Fatality cost	10 million euro	2 million euro
Lethal radiological dose	0.2 Sv	7 Sv
Major injury cost	3 million euro	200,000 euro
Core damage cost	15 million euro	460 million euro
Regulator inspection cost	3 million euro	2.6 million euro

3.8 Summary

We have now outlined the decision *modeling* for the incident management prototype. It is the goal of the incident advisor prototype to select a preferential decision alternative from the list of potential options and provide a technical justification for the basis of the decision. The primary focus of the decision model utilizes the structure drawn from influence diagrams. Within this decision model, we identified of six major parts:

- Decision alternatives – these include the options specific to the incident.
- Incident specific elements – these include the possibility for repair, the type of failure mechanisms present, and other unique features related to the incident.
- Boundary conditions – these include the plant state and time until the next outage.
- Plant upsets – these include initiators such as transients/leaks that lead to complications.
- Plant response – these include the plant system response to any upset conditions.
- Outcomes – these include the outcomes of interest to the decision maker.

Further, we selected an additive version of the multi-attribute PI for use in ordering decision options. The form of the PI was found to be:

$$\begin{aligned} PI &= w_{economics}\, u(economics) + w_{dose}\, u(dose) + w_{accidents}\, u(acc.) \\ &\quad + w_{safety}\, u(safety) + w_{stakeholders}\, u(stakeholders) \\ &= PI_{economics} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{stakeholders} \end{aligned} \tag{34}$$

Since we utilize disutilities, we seek to minimize the PI for the selection of a decision. As noted in the PI equation, we identified (with our decision makers) a total of five performance measures and their associated weights.

Note that the PI shown is only a point estimate. This value is only used to give the user a "feel" for the decision alternative outcomes. For decision making, one must rely on the expected value of the PI. Consequently, the prototype performs Monte Carlo sampling on the input variables to PI as part of the expectation calculation. But, once the PIs are known for the decisions, we rely on the decision rule to select a preferential decision. In our case, we use disutility functions, which implies that we desire to minimize negative outcomes. Thus, decisionS with low PI are preferred, or:

$$Decision(1) = \min(E[PI]_i) \tag{35}$$

where *Decision*(1) is a the preferred decision alternative and $E[PI]_i$ is expected value of the PI for the *i'th* decision alternative.

We used AHP to determine the value tree performance measure weights. When we use AHP to determine the value tree weights, we are comparing one entity (the worth of cost) against another entity (the worth of worker safety). These value tree measures are within an order of magnitude from one another (one is not 100 or 1,000 times more important than the other in the context of incident management). Further, they do not have arbitrary outcomes (for example, on the cost measure, the focus is "cost" as an impact, not on any one particular value for cost). These conditions on the weights are consistent with those suggested by the developer of AHP (Saaty, 1980; Saaty, 1997).

Following the performance measure weight determination, we proceeded to determine the disutility functions for each performance measures. We initially relied upon AHP to derive the disutility functions. But, we found several limitations in the use of AHP for disutility elicitation:

1. AHP yields disutilities that imply the decision makers are very risk prone.

2. AHP preserves the certainty equivalent only for the low scale regions (e.g., zero to one million euro). In higher regions, the certainty equivalent is suspect. For

example, the AHP results suggest that the decision maker would only be willing to pay 20,000 euro to avoid a "50-50" lottery (say a flip of a coin) where the two outcomes are either lose nothing or lose 1 billion euro. In reality, the decision makers would be willing to pay much more than 20,000 euro to avoid large losses of this type.

3. AHP disutility curves are extremely sensitive to small changes in the initial scale regions. In other words, it is possible to obtain dramatically different disutility functions depending on the nuances of setting up the AHP scales.

4. The sanity checks performed on the AHP results showed several inconsistencies. For example, the AHP application indicated that a lethal dose should be approximately 0.2 Sv, when in fact it is approximately 7 Sv.

If we return to the PI equation, it is evident that any inconsistencies must be resolved via modifications to either the equation itself or the parameters of the equation. Specifically, we indicated that three possible changes could be made:

1. Modify the performance measure weights
2. Modify the functional form of the performance index equation (e.g., move to a non-linear form)
3. Modify the disutility functions themselves

Only the last item (3) show promise in relieving the inconsistencies that were found in our modeling of the disutilities and performance measure weights. From that notion, we developed an approach called "measurable equivalence" in order to construct the required disutility functions. This approach allows us to ensure that the maximum outcome (the worst case) for each performance measure has about the same level of "consequence." Further, we utilize performance measure indifference points in order to constrain the disutility function, where the constraint is by actual measurable equalities. This second feature is used to bring real data into the decision process while simultaneously helping to reduce the subjectivity present when utilizing preference information. In order to transpose equivalencies from one performance measure to another, we needed to have at least one disutility function fully specified. We chose to determine the cost disutility, where we determined a set of points on this function by way of lottery equivalence questions poised to our decision makers. We then proceeded to determine the remaining disutility functions by applying the measurable equivalence approach, which calculates disutility points via the equivalence weights and disutilities. The last step in this process is deliberation with the decision makers where adjustments to the disutility functions are considered.

The decision model embodied in the prototype is a generalized influence diagram. Other analysis modules support the main decision model, including economic impacts, risk analysis, and worker safety. Interacting with these modules is a knowledge base that supports the generation of an incident-specific decision model. But, once the decision model is constructed, we still need to analyze the model. It is the focus of the next section to outline the subsequent analysis of the decision model.

" The only law is that there is no law." — John Archibald Wheeler

4 Analysis Heuristics for Incident Decision Making

It is the goal of the incident advisor prototype to select a preferential decision alternative from the spectrum of available options and provide technical justification for the basis of the decision. In this section, we outline the heuristics used to analyze the decision model that was outlined in the previous section. For the analysis discussion, we separate the focus into three subsections, (1) treatment of uncertainty in the decision process, (2) treatment of deterministic aspects of the analysis, and (3) treatment of aleatory aspects of the analysis.

As indicated earlier, the general focus for modeling in the incident management framework is through the use of influence diagrams and related models. While the influence diagrams allow for a depiction of the "big picture" of the decision problem, one still needs to have rules to solve the model in order to determine preferential decision alternatives. As we will discuss, potential analysis solution approaches are available via decision trees, Markov models, or event simulation.

Incident management begins with the realization that an incident has occurred that requires attention. From the incident impact point, one represents its implications via a model like the influence diagram. Using supporting modules such as the value tree, disutilities, the PRA, and economic impacts, the decision model will be solved to provide decision advice. This general framework for the incident decision making is shown in Figure 37, where we single out simulation as the analysis method.

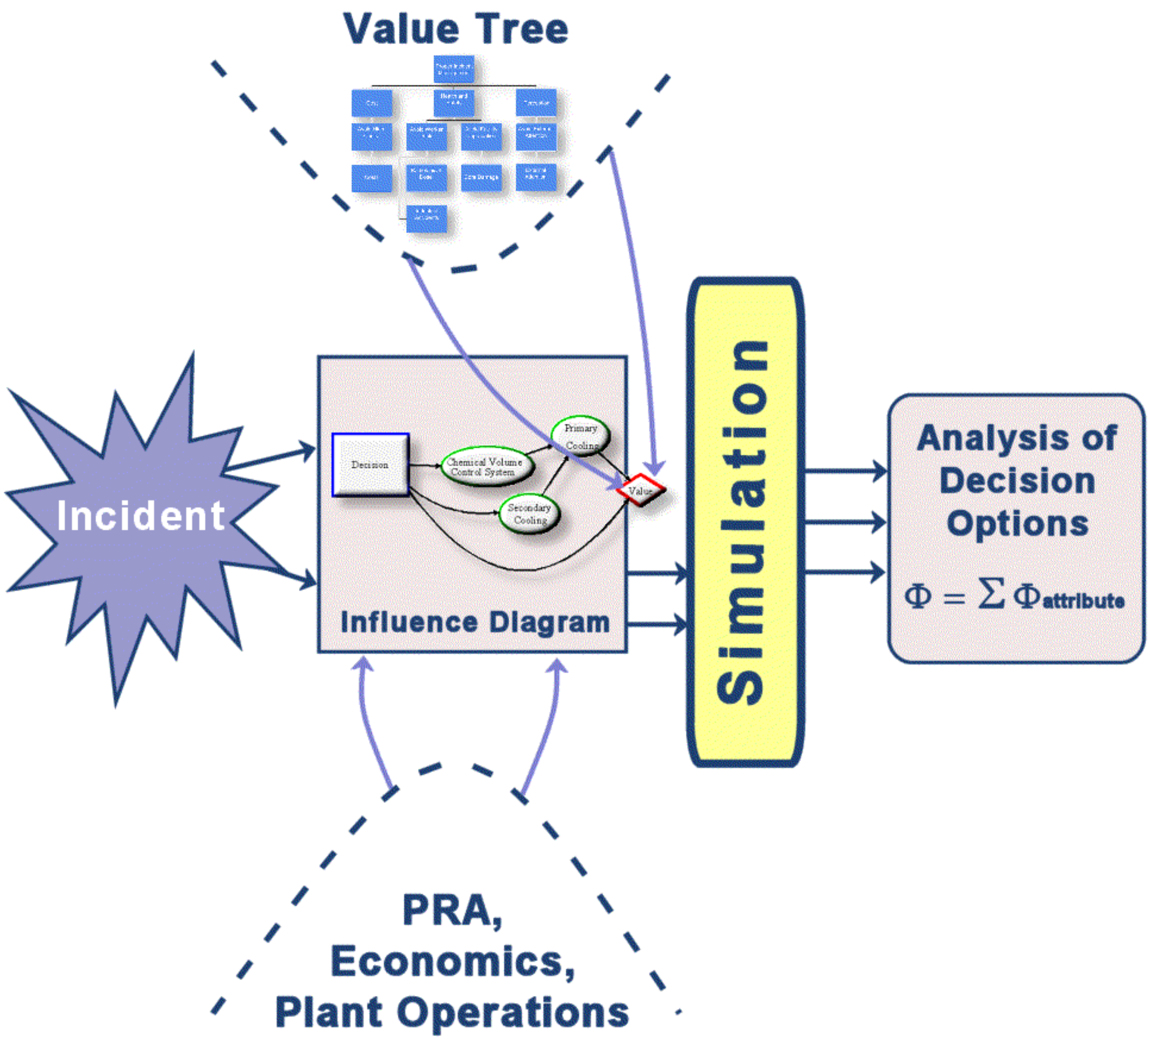

Figure 37. Overview of the decision analysis process where simulation is used to solve the influence diagram decision model.

4.1 Treatment of Uncertainty in the Decision Process

The analysis of the incident management decision model involves the determination of preferential decisions based upon expected disutility. The expected disutility for our prototype has been defined as a function of five performance measures, as indicated in the performance index for one decision alternative:

$$\begin{aligned} PI &= w_{economics}\ u(economics) + w_{dose}\ u(dose) + w_{accidents}\ u(acc.) \\ &\quad + w_{safety}\ u(safety) + w_{stakeholders}\ u(stakeholders) \\ &= PI_{economics} + PI_{dose} + PI_{accidents} + PI_{safety} + PI_{stakeholders} \end{aligned} \tag{36}$$

In order for us to determine the expected value of this equation, we must integrate over the uncertainties present in the performance measures and associated weights for a decision alternative:

$$E[PI] = \int_0^{\infty} \left(\sum_{i=1}^{5} w_i\, u(x_i \,|\, y_i) \right) \pi(\overline{\mathbf{x}}_i \,|\, \overline{\mathbf{y}}_i)\ d\overline{\mathbf{x}}_i \tag{37}$$

where i indicates the performance measure; w_i is the weight of the measure; u_i is the disutility for the measure; π( ) is the epistemic uncertainty over the decision analysis parameters and models; x is the decision analysis parameters and models (which, in general, is a vector of factors for each decision alternative and performance measure); and y is the boundary conditions of the problem (i.e., the evidence or the facts). Since we are dealing with disutilities for the prototype system, we seek to have decision options with low E[PI]. It is evident from Equation 37 that treatment of uncertainty is not only critical but also required for determination of the decision rule based upon expected values. While point estimate calculations may be useful in a qualitative fashion, they are not supported in the literature as the basis for decision making. Consequently, we discuss the relevant aspects of the uncertainty analysis prior to presenting details of the supporting (e.g., PRA, economics) analysis.

4.1.1 Uncertainty Concepts

Our decision framework represents what the decision makers think (via decision options, chance events, and related outcomes) and how they feel (via preferences) about the reality behind a specific incident. As such, this framework and its attendant models embodies our “model of the world” (Winkler, 1972; Apostolakis, 1995). Within our model of the world, there are two basic types of reality abstractions, aleatory and deterministic. While most individuals are familiar with the later, let us discuss both since they are very relevant to our overall decision framework.

Aleatory models represent randomness (or a stochastic process) in an outcome of interest. For example, flipping a coin is an aleatory process and is frequently modeled by using a binomial distribution to characterize the number of heads (or tails) that we see for a given number of flips. Here, the random, but observable, quantity is the number of heads (or tails), where the key here is observable outcomes. Since "probabilities" are not observable quantities, we do not have an aleatory model directly for probabilities. Instead, we rely on models (such as a binomial) to estimate probabilities for certain outcomes (e.g., two heads out of three tosses of the coin).

In the PRA, we utilize numerous aleatory models, generally to represent component failures and, to a lesser degree, operator actions. When modeling component failures, we estimate failure probabilities by utilizing observable such as time to failure (for a Poisson type failure) or the number of failures (for a binomial type failure). Here, the random, observable parameters are time or number of failures.

If time is the observable quantity of interest (specifically, the time until a failure), we can not say for certain how long this time will be for any one component. Consequently, the time to failure exhibits stochastic behavior. In most reliability modeling, this stochastic behavior is represented using a Poisson model. Further, a “rate of failure” is typically estimated for this process, but this rate is not a measurable quantity (it may be inferred though). If we assume that this failure rate is constant, then we can utilize the Poisson model to estimate a failure probability via the equation: probability = 1 - EXP(- λ T), where λ is the constant failure rate and T is the time period of interest. One may ask why should we be concerned with having an observable quantity up front (e.g., time to failure) when we still end with an equation in terms of a probability. The reason for this distinction is that we could evaluate the model in two ways – we could have probability as the end metric or time as the end metric. For example, we could structure a reliability model such that the result would be expressed as a "mean time to failure." But traditionally, reliability models are generally used in terms of probabilities.

So, the equation P = 1 - EXP(- λ T) is a model representing the stochastic nature of our modeling situation. In other words, the equation is an aleatory model.

A deterministic model is quite different than an aleatory model. This model type represents situations where the observable quantity will be known (not unknown as in the aleatory case) given a certain set of parameter values. For example, the equation $E = mc^2$ is a deterministic model. Here, if we know the mass m and the speed of light c we presume to know the resultant energy. Now, we may not know the energy precisely, and in fact we will know it up to the uncertainty in the model parameters (m and c), but what we can predict is that the energy will not be stochastic like for the case of time to failure for a component. Another example of a deterministic model is for the situation shown in Figure 38. If we wanted to model block movement down an incline, we would need to know at what time the block will pass point A after it passes point B (assume it is already in motion at point B so we do not have to worry about the initial movement). One might question why do we not ask, in our model, for the velocity of the block at point B. If we know the velocity and we know the distance between A and B, we can calculate the time. But, velocity is not an observable quantity. Observable quantities in this case are the block position and time. From these two quantities, we can obtain the velocity since we know that velocity = $\partial x/\partial t$. But, this equation is deterministic, not aleatory – there are no probabilities involved.

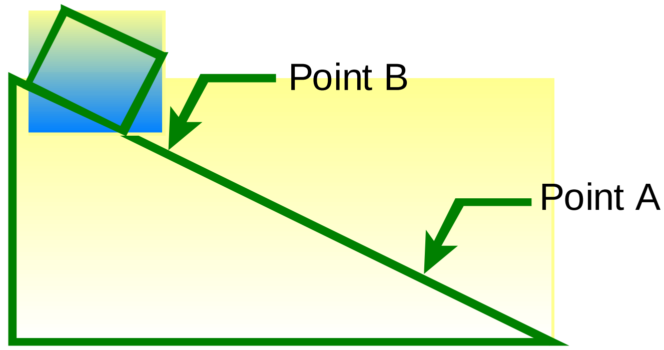

Figure 38. An example from physics of a deterministic model.

Each of the two model types (aleatory and deterministic) has parameters that may have uncertainty. Consequently, we introduce the notion of epistemic uncertainty. Epistemic uncertainty represents how accurate our state of knowledge is about the model, regardless of the type of model (Eerola, 1994). Within this uncertainty we include the impreciseness on the model parameters (i.e., parametric uncertainty), issues on the model itself (i.e., modeling uncertainty), errors during quantification, etc. While the parametric uncertainty is generally the one type of epistemic uncertainty that does get evaluated (if any), it is possible that other uncertainties are more important. As noted by Ang and Tang (1975), "the uncertainties associated with such predictions or model error may be much more significant than those associated with the inherent variabilities."

If we are dealing with an aleatory model [e.g., $P = 1 - EXP(- \lambda T)$] or a deterministic model (e.g., $V = \partial x/\partial t$) and if any parameter of these models is uncertain, then the model has epistemic uncertainty. Stated another way, if the inputs to a model have epistemic uncertainty, the output of the model has epistemic uncertainty. Back to our two models, if the failure rate λ is uncertain (which it always is), then our probability P is uncertain and if the block position x is uncertain (which it always is), then our velocity V is uncertain. Since most, if not all, of our engineering models rely on observable quantities, our models have epistemic uncertainty since:

1. We can not be absolutely precise when we measure parameters like position and time.
2. We can not have an infinite amount of measurements on the observable quantities.
3. The observable quantities exhibit stochastic behavior.

Now that we have defined the three key uncertainty terms, aleatory, deterministic, and epistemic, we can dissect the decision analysis framework to see where key uncertainties lie and methods for their treatment.

4.1.2 Uncertainties within the Decision Analysis Framework

The decision analysis framework encompasses a variety of models, both aleatory and deterministic. For example, since the nuclear power plant PRA is an integral part of the framework, its associated models plays an important part. Included within the PRA are aleatory models (probability that a component fails to start, probability that a component fails to operate, probability that an operator fails to restore a component, etc.) and deterministic models (the system fault trees, the accident sequence event trees). But, the overall framework entails more than just the safety model. As we show in Figure 37, several important modules, in addition to the PRA, are required. In that figure, we noted two critical deterministic models, the value tree and the decision model (e.g., an influence diagram). Also, the resultant metric out of the overall framework, that given by the PI equation, is a deterministic model.

Additional detail on the decision framework models is shown in Figure 39. Here, we note that there are three primary analysis modules supporting the decision model: (1) economics, (2) safety analysis, and (3) worker actions. These three areas affect different parts of the decision model, specifically within the context of the value tree. An influence implies that if an impact is identified that does not affect one of the five value tree attributes, then it is not relevant to the quantitative portion of decision process.

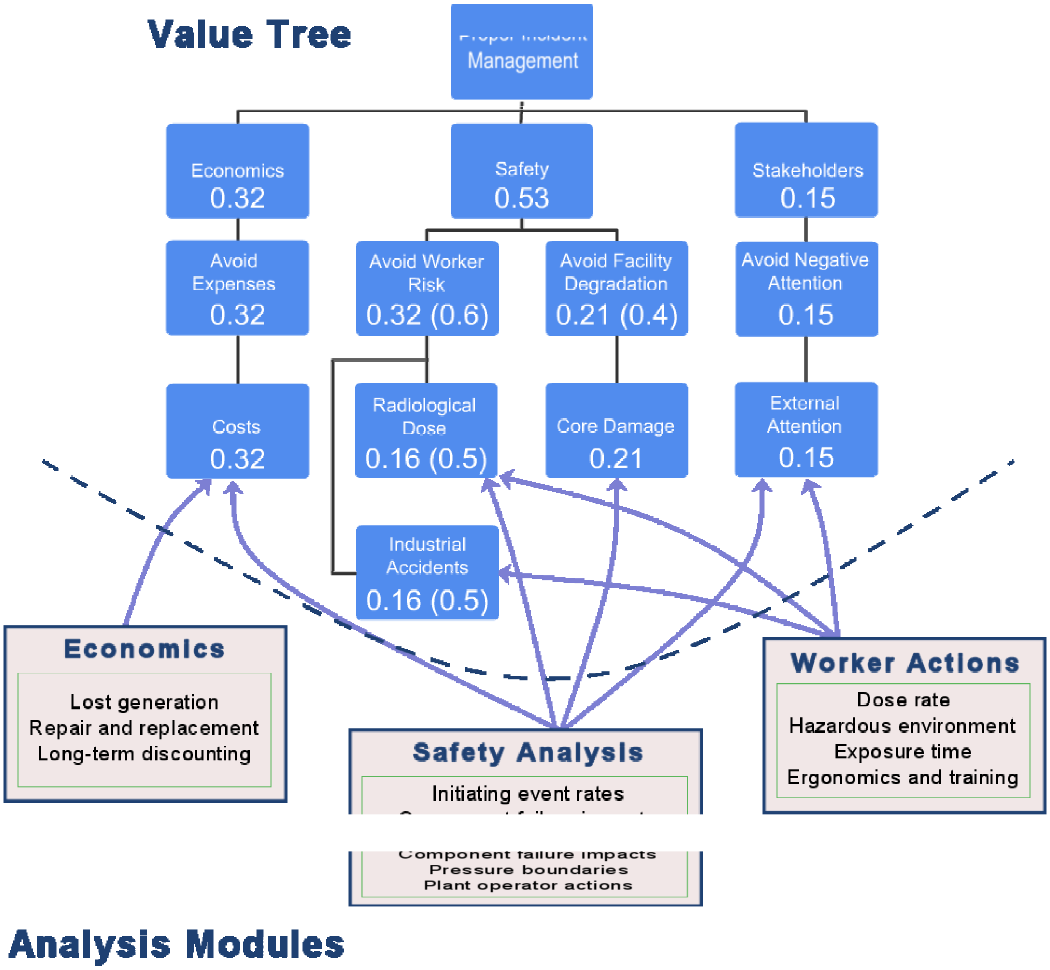

Figure 39. Decision making framework modules and their influence on the value tree.

The majority of the analysis model types listed in Figure 39 are deterministic, with the only aleatory models appearing under the safety analysis portion of the figure. While most of the modeling here is deterministic, all of the analysis modules do have epistemic uncertainties that must be factored into the calculation of expected disutility. For example, under economics, the cost of lost power generation is not a fixed monetary value – the cost of power varies from plant-to-plant and as a function of time. This imprecision in the value used for the replacement power costs will need to be factored into the calculation any place that this cost variable is used in the decision model as indicated by Equation 37.

The overall decision framework can be further subdivided as shown in Figure 40. Here, we identify the primary parts of the analysis in five categories: utility theory, plant operation, engineering judgement, safety assessment, and economics. For each of these categories, we provide examples of either deterministic or aleatory models (as applicable) and related epistemic uncertainties that may be found within the category. We also identify how these categories might affect the components of the decision model through impacts to the decision alternatives, the chance events, or the outcomes. We will discuss unique facets of uncertainty for these categories in turn.

4.1.2.1 Uncertainties in the Value Tree and Disutilities

Models such as the value tree, AHP, and measurable equivalence are deterministic. While deterministic, each has parameters or issues that give rise to epistemic uncertainties. It is unfortunate though that these uncertainties have received sparse attention from decision analysis researchers. Some work has been conducted on the model uncertainty portion of epistemic uncertainty, where the consistency between preference elicitation methods was compared (Pöyhönen and Hämäläinen, 2001). In addition, much of the prior work evaluating the parametric uncertainty has looked at either judgement on the preference intervals or the direct assignment of probability distributions (Escobar and Moreno-Jiménez, 2000). Some of the work related to the assignment of probability distributions is limited for our application due to the assumption that the AHP matrix entries are independent. Further, we question the assignment of probability distributions on the AHP matrix entries since this practice presumes that an entry is a random or stochastic variable. If an AHP matrix entry is indeed a random variable, is *could* be represented by an aleatory model. But, aleatory models are tied to an observable quantity that displays stochastic behavior (such as the time to failure for a component). In the case of an AHP matrix entry, the entry itself represents the decision maker's preference. As such, the availability of an associated stochastic, observable quantity is illusive.

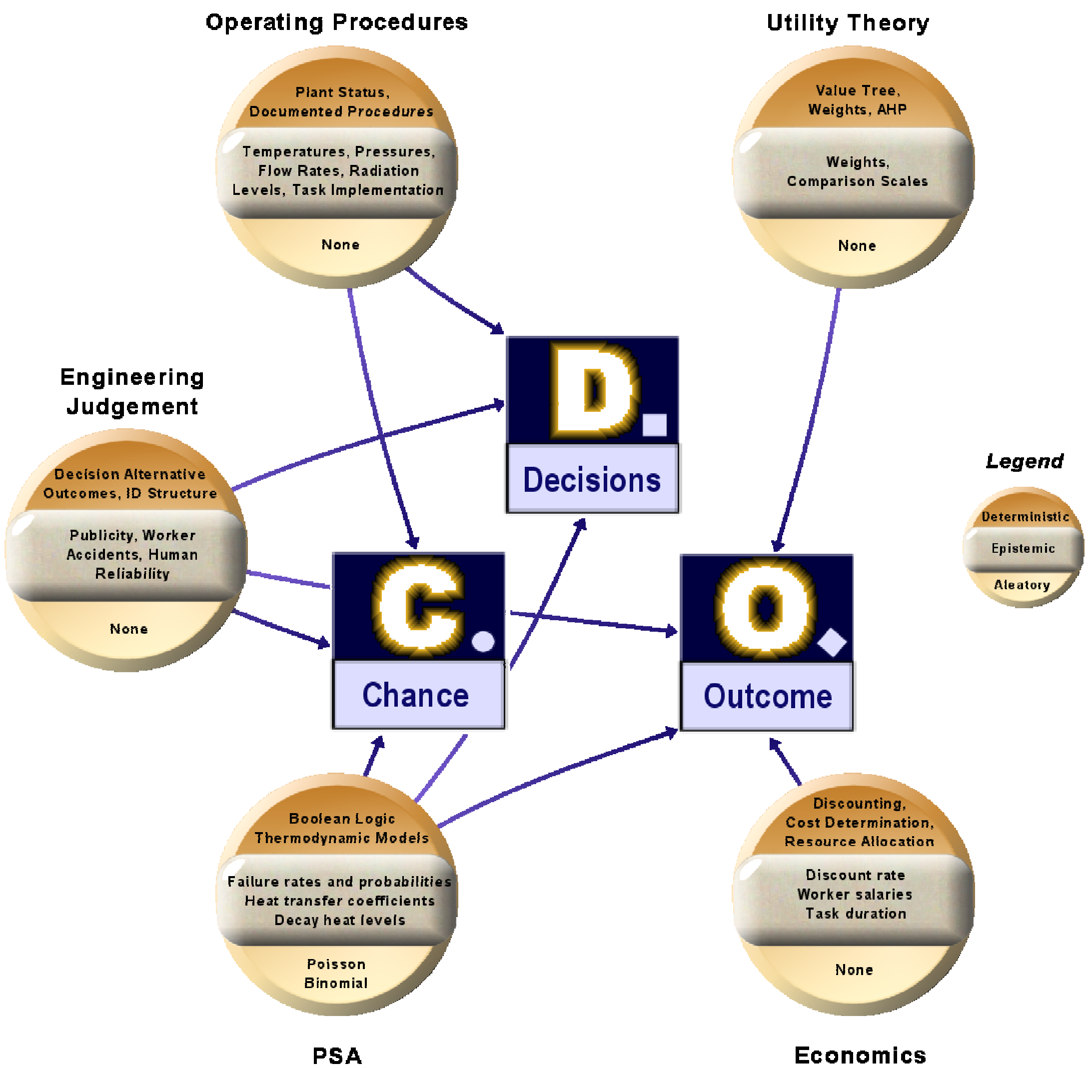

Figure 40. The decision making framework identifying deterministic and aleatory models, epistemic uncertainties, and their influence on the decision model decision alternatives, chance events, and outcomes.

Instead, it is our belief that the variability from the AHP process is entirely due to epistemic uncertainty, namely that related to model uncertainty where the decision maker is not entirely sure of the numeric AHP scale value for a particular pairwise comparison.

Since AHP is used often in decision making arenas and was used for our determination of value tree weights, we decided to investigate the variability in AHP-based results. To collect AHP information, we provided comparison worksheets to our group of decision makers. An example of the worksheet used during the exercise is shown in Figure 41. In this figure, we ask the decision maker to perform a pairwise comparison of the cost attribute for four outcomes: (1) no cost, (2) 5,000 euro, (3) 50,000 euro, and (4) 500,000 euro. The ranking used is the typical one-to-nine AHP scale, where a scale of "1" indicates two outcomes are equally preferred (indifferent) while a "9" indicates that one outcome is absolutely preferred over another outcome. It is in this ranking process where we need to question and quantify the uncertainty on disutility in order to correctly reflect the uncertainty embodied in the overall decision process.

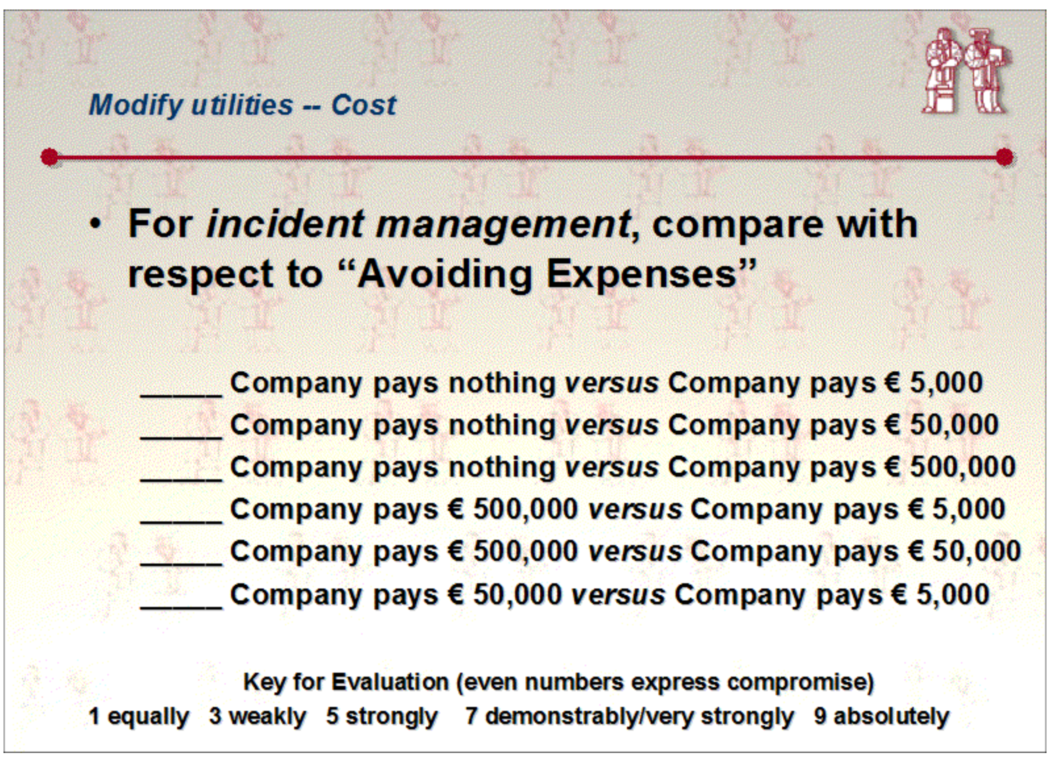

Figure 41. Example of the disutility determination AHP worksheet for the cost attribute.

The scales assigned during the AHP process are representations of preference. As such, the decision maker is given time to reflect on the pairwise comparisons. The judgement here is not akin to decision making under time duress or other pressures. Discussion and deliberation is an important part of the process, not only to help the decision makers understand the process itself, but to provide reference points related to specific attributes (e.g., a cost of one million euro is approximately one lost day of production at a nuclear power plant). Given the conditions under which this preference information is obtained, it was our belief (read assumption) that for the types of attributes indicative of incident management, the decision maker is able to specify his or her preference to an accuracy of the scale, plus or minus one integer value away from this scale value. Thus, if the decision maker indicates that for the pairwise comparison of two outcomes (say paying nothing versus paying 50,000 euro) the scale selected has a value of “4” then there is some chance the decision maker really wanted either a “3” or a “5.” Further, we assume that the decision maker *would not* want to assign a scale of “2” or ”6” in this case. Note that we are not going to attempt to use fractional scale values (e.g., 3.2, 4.5) with AHP even though some researchers have used such approaches. If the most accurate a decision maker can be is approximately one scale unit, then attempting to justify subdivisions to the scale would not be supported by the epistemic uncertainty. Further, attempting to educate the decision maker on the subtleties between a scale value of 3.8 and 4.6 (“sort of weakly?” versus “almost strongly?”) would only hinder the preference elicitation.

Assuming that the AHP scale accuracy is to one scale unit is just that, an assumption. Thus, we need to look for evidence that would either support or discount this assumption. We found this supporting evidence in two places, one qualitative and one quantitative.

First, during the disutility workshop, we noticed a couple of times that the decision maker would record a scale value, then change that value, presumably after further thinking about the pairwise comparison. In all cases when this happened, the change was at most one scale unit (either up or down). It could have been possible that the decision maker had made notes (either mentally or in written form) not on the worksheet that would indicate larger changes to the scale values. But, we did not notice any evidence of this.

Second, following the workshop and the numerical disutility evaluation, we took all four sets of worksheets to compare how one decision maker specified the AHP scale value against the others. During the workshop, there was generally discussion prior to the start of the worksheet tasks to ensure everyone understood the tasks. But, once the pairwise comparisons started, in general each (of the four) decision makers operated independently. Note though that the workshop took place in a single room without physical barriers preventing communication, nor did we discourage discussion. But, we made sure prior to filling out the worksheets that the decision makers understood that we would provide a forum for deliberation later to ensure consistency between the decision makers. Thus, during the initial AHP application, the decision makers were free to express their own thoughts.

The results from the disutility workshop were pairwise AHP comparisons for the four decision makers. Consequently, it was realized that we could query these four sets of data to see how much did the AHP scale values change from one decision maker to the next. Then, we can utilize this comparison to bound the AHP scale accuracy since an individual decision maker should vary less than different decision makers working independently (again, an assumption). The scale data analysis took the pairwise comparisons for five worksheets (for four decision makers, giving a total of twenty worksheets) where the specific worksheet elicitation is described in Table 16. We then calculated the standard deviation between the recorded AHP scale for each comparison. The results of this data analysis, looking at the variability between decision makers, are shown in Figure 42.

Table 16. Information for five of the AHP disutility worksheets.

Worksheet Identifier	Attribute	Outcome Range
Cost #1	Cost	0 to 500,000 euro
Cost #2	Cost	1,000,000 to 1,000,000,000 euro
Cost #3	Cost	100,000 to 100,000,000 euro
Dose #1	Radiological dose	0 to 0.1 Sv
Dose #2	Radiological dose	0 to 100 Sv

As we can see in Figure 42, the variation as measured by the standard deviation ranges from zero (all four decision makers indicated the same AHP scale value for the particular comparison) to two. The average variation for the five worksheets was found to be 1.2. When we evaluated the disutility cases for industrial accidents and precursors to core damage, we found the average variation was about 1.7, but these two cases had five scale points (e.g., for industrial accidents, we had none, minor injury, major injury, one fatality, and 20 fatalities as the outcomes) rather than the four points for the cost and dose worksheets.

It is interesting to note that for the case of the cost worksheets (Cost #1, Cost #2, and Cost #3) we did not see a deviation between the sheet variability (each averaged about a value of 1.4), even though the cost range represented drastically different outcomes. The two radiological dose cases did show slight changes in the variability where the "lesser impact" case (Dose #1) had an average variability of 0.8 while the second case (Dose #2) had an average variability of 1.1.

We conclude, based upon both the qualitative and quantitative insights, that the application of AHP for the performance weight yields a process that is able to specify preference to an accuracy of the scale, plus or minus one integer value away from the scale value.   Since deliberation was used in determining the final performance weight results, we need to be able to specify an uncertainty on the weights based upon the potential variation. In other words, we utilize the final performance measure weights as appropriate expected value, but will then apply an uncertainty one could expect to see out of the direct application of AHP.

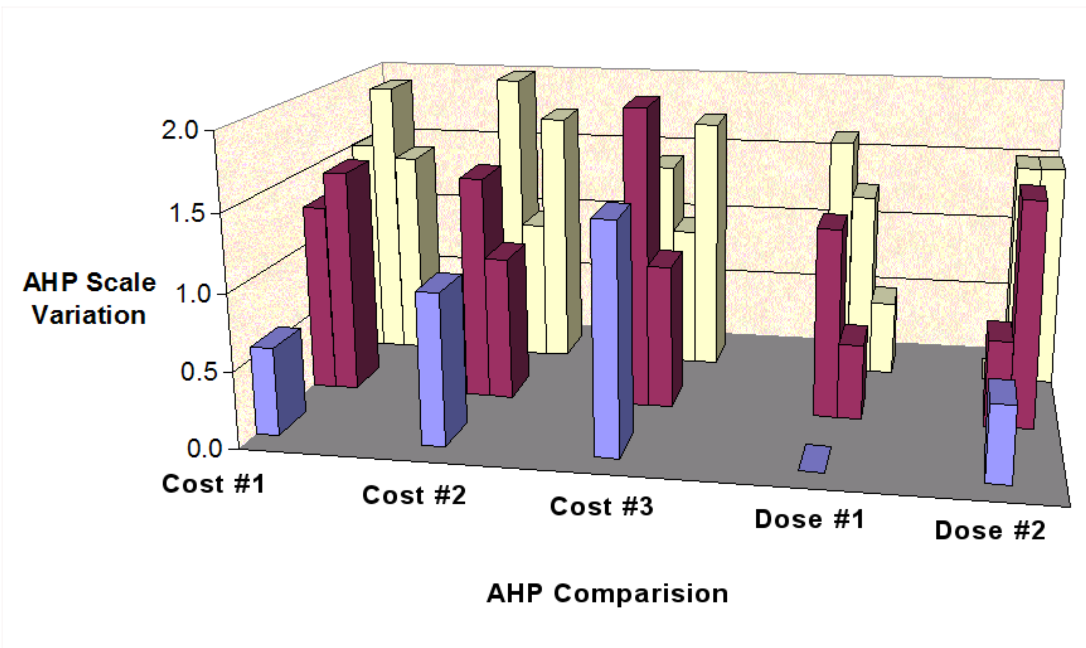

Figure 42. Variation in AHP scales assignments for cost and radiological dose evaluations for four independent decision makers. (The front-most bars represent the variation in the AHP first column, the middle bars represent the variation in the AHP second column, and the rear bars represent the variation in the AHP third column.)

To obtain insights into the uncertainty on the variability one might find when applying AHP, we developed synthetic AHP matricies for a four-point (e.g., scale category I, II, III, and IV) comparison. Here, the synthetic matrix of scale N is defined as:

$$
\begin{array}{c|cccc|}
 & \text{I} & \text{II} & \text{III} & \text{IV} \\
\text{I} & 1 & N & N & N \\
\text{II} & 1/N & 1 & N & N \\
\text{III} & 1/N & 1/N & 1 & N \\
\text{IV} & 1/N & 1/N & 1/N & 1
\end{array}
$$

We then changed the AHP scale values from N=1 to N=9 and plotted the resultant disutility curves. The results of this calculation (for select values of N) are shown in Figure 43. From this figure, one gets a sense of the type of variability we might expect to see in the applying AHP – variability that will manifest itself as part of the epistemic uncertainty in the overall decision process. We need to note a couple of items here:

1. Since the disutility is constrained between zero and unity, there is no variability exhibited by the best or worst outcome disutility values. By definition, the best outcome has disutility of zero while the worst outcome has disutility of one.

2. The variability in the disutility is relatively small. For example, jumping from the N=2 to the N=3 disutility for scale category III only increases the disutility value by about 6%. For the same category, going from N=7 to N=8 increases the disutility by less than 1%.

3. As N increases, the percent increase in the disutility value decreases. Thus, there is an inverse relationship between the AHP weight and the potential variability. This relationship is plotted in Figure 44. The inverse relationship between variability and AHP scale values – which is manifested in the epistemic uncertainty on AHP – has been noted qualitatively by others. For example Accorsi, Apostolakis, and Zio noted a "saturation of the (AHP) scales" problem during their elicitation of decision maker preferences. (Accorsi, Apostolakis, and Zio, 1999)

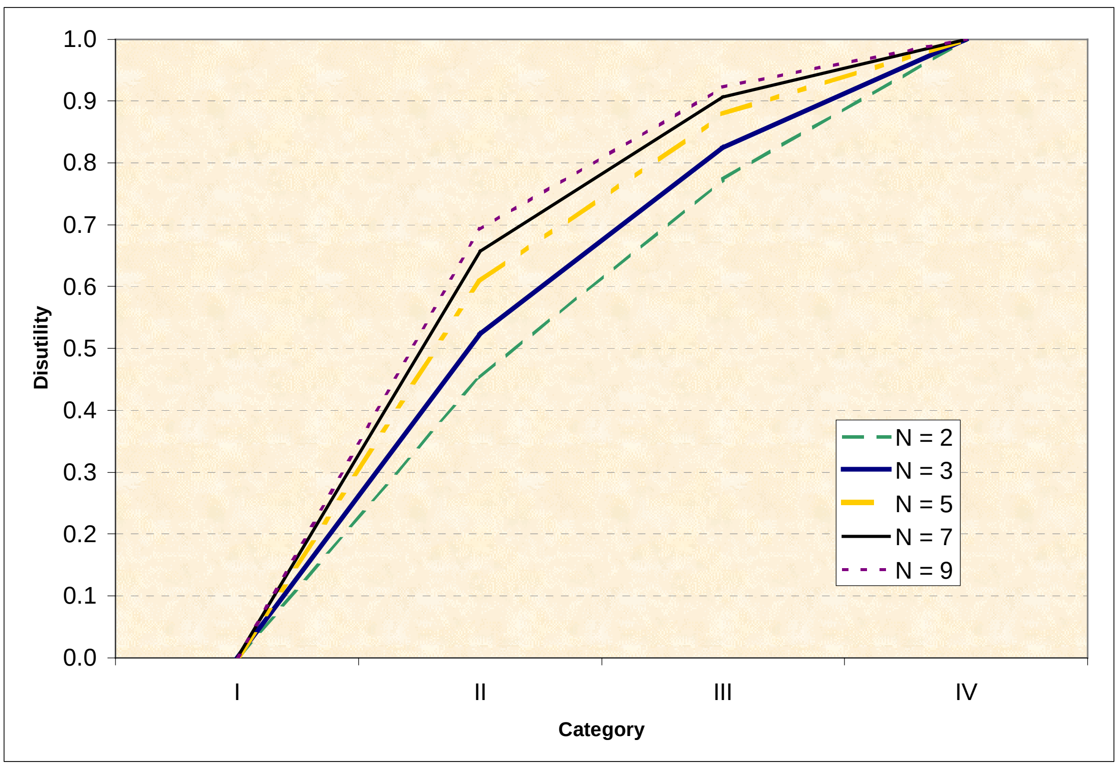

Figure 43. Disutility variability for a four-point curve where the AHP weights are modified from N=2 to N=9 for a synthetic matrix.

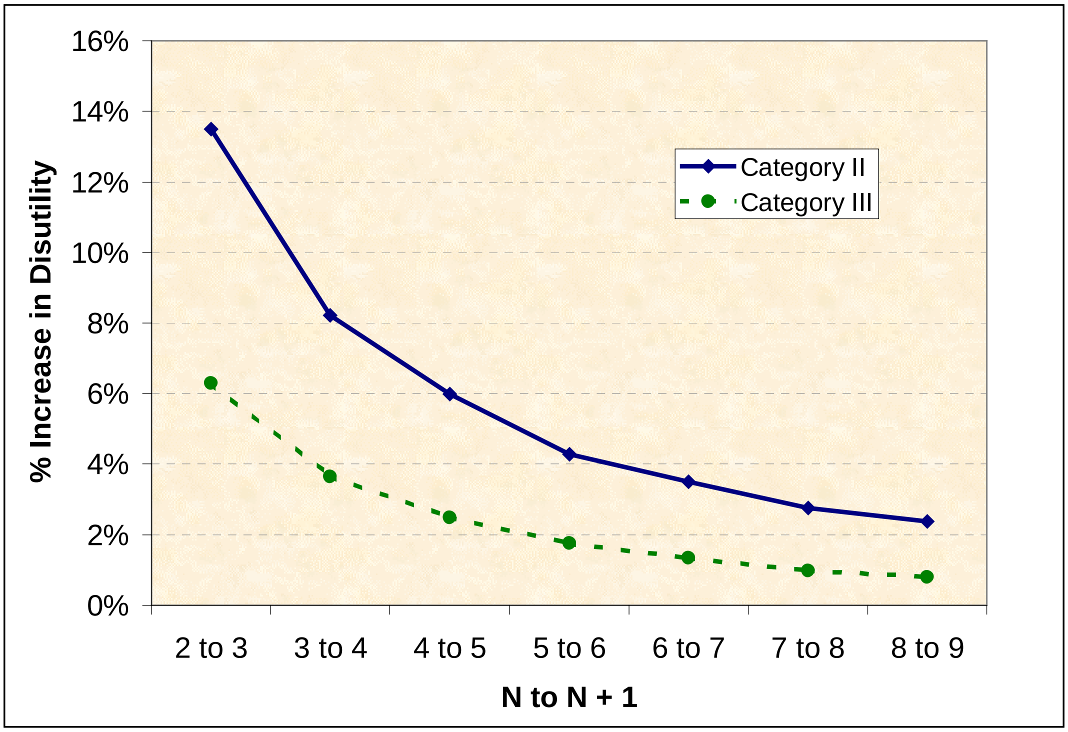

Figure 44. Variation in disutility values as a function of change in AHP scale.

Ultimately, as part of the overall decision framework, we can use the information shown in Figure 43 and Figure 44 to approximate the epistemic uncertainty related to the use of AHP (recall that AHP was used to determine the performance measure weights).

Since the overall variability seems to be relatively small (based upon the two previous figures), we are going to simplify the Monte Carlo uncertainty sampling by forcing the upper and lower bound to represent a uniform distribution on potential values. Further, we selected a representative value of 5% as representing the uncertainty. A slightly more sophisticated approach would be to utilize a triangular distribution where the midpoint is defined by knowing the mean and the upper and lower bounds. But, we do not expect the numerical results to differ much between these two approaches, consequently we opt for the routine (the uniform distribution) that both (1) minimizes the software development time and (2) takes less computation time during the numerical evaluation.

4.1.2.2 Uncertainties in Plant Operation

In theory, our decision analysis model of the world relies on the knowledge of physical attributes of the nuclear power plant such pressures and temperatures. Information related to these physical attributes impact calculations throughout the decision model, including success criteria in the PRA, selection of decision alternatives, and possible outcomes. But, for the analysis described in this report and for the decision advisor prototype under development, we do not decompose the decision model down to the "thermodynamics" level. While there certainly are epistemic uncertainties related to the parameters and models for plant operation, they simply were outside the scope of the present analysis.

Note though that some operational parameters do creep into the analysis, and when encountered, may be subjected to sensitivity analyses. For example, in Figure 45, we illustrate the sensitivity calculation for decision preference where (1) the time to scheduled shutdown and (2) the leak rate for a faulty steam generator tube are allowed to vary. As the physical plant parameters are varied, notice that the preferential decision alternative (isolate steam generator, shutdown the plant, continue operation with normal water makeup) changes. At the point of the decision, both the time until the next schedule shutdown and the leak rate will be known by the decision maker – hence the desire to treat these parameters via sensitivity calculations rather than as part of the model epistemic uncertainty.

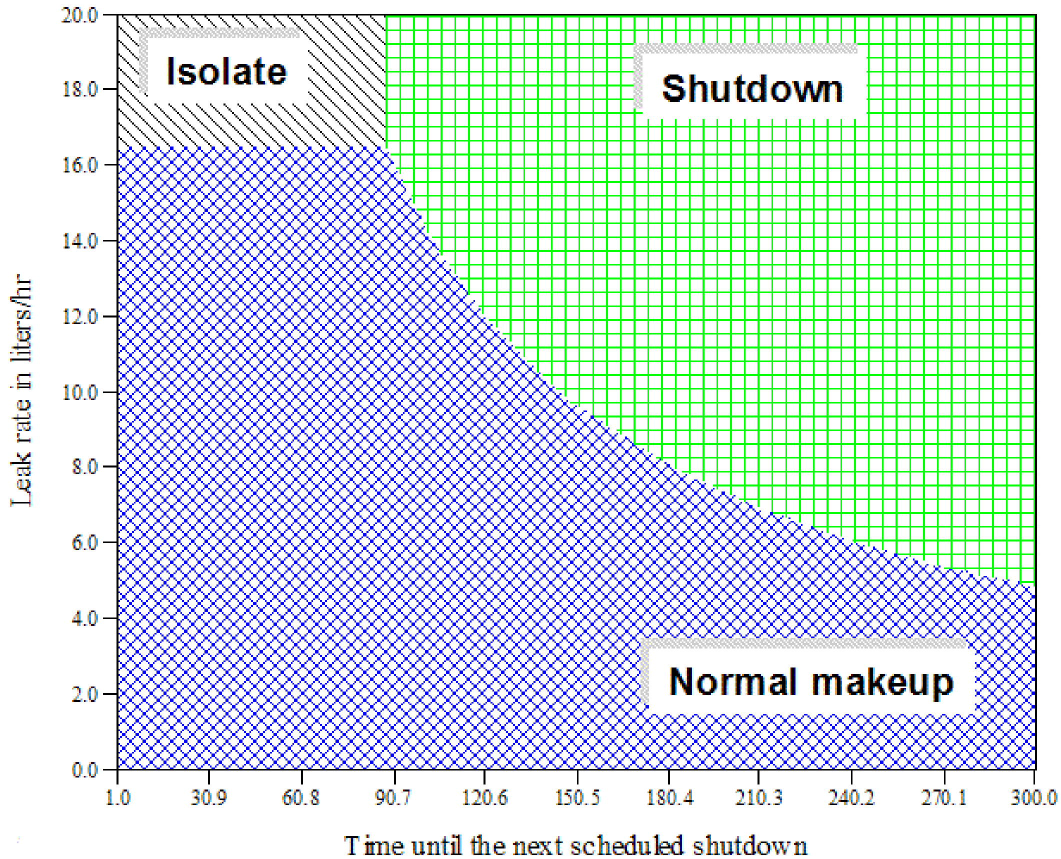

Figure 45. Example of the sensitivity in decision alternative preference as a function physical plant parameters (1) time to scheduled shutdown and (2) leak rate for a faulty steam generator tube. The shading indicates the preferential decision alternative, either normal primary makeup, isolating the leak, or shutting the plant down to repair the leak.

4.1.2.3 Uncertainty in Engineering Judgement

Judgement plays an important role in many engineering calculations simply due to the fact that we can not model every aspect of a problem nor, for those things we do model, can we represent detail to a microscopic level. Instead, engineers tend to focus on what is thought to be important and then model details to a macroscopic degree.

In our work detailed in this report, judgement on outcomes such as the degree of negative attention (or "bad publicity") or the potential for worker injuries or fatalities has an important role in the decision analysis. Consequently, tied to each potential decision alternative is the necessity to indicate the level of outcome for the value tree attributes "stakeholder attention," "radiological dose," and "industrial accidents." The uncertainties on these attributes could be evaluated through expert opinion like we did for the disutility determination. Two of the attributes (stakeholder attention and industrial accidents) consist of disutilities measured on an ordinal scale. The discrete nature of these attributes reduces the need for formal uncertainty determination of their outcome (for example, we can not have ½ a fatality). Nonetheless, the decision prototype should use epistemic uncertainties on the attribute outcomes similar to the approach we used for disutilities.

A more formal analysis of the uncertainty on factors residing from expert judgement could utilize Bayesian noninformative uncertainty distributions. One very useful distribution that has been developed is the constrained noninformative, where the constraint is that one knows the mean value of the factor (Atwood, 1996). For many situations of expert judgement, we are fortunate to know the parameter mean value – thus this noninformative distribution is a convenient match to the realized intent of the judgement. Use of this distribution results in either a gamma or beta distribution (depending on the parameter being modeled), where the indicated mean value matches the mean value of the gamma/beta distribution. Further, this distribution is an extension of the Jeffreys prior that plays a central role in much of noninformative statistics (Martz and Waller, 1991). But, the constrained noninformative distribution utilizes the user-specific mean value as part of the prior distribution, whereas the standard application of a Jeffreys prior utilizes the mean value of the Jeffreys distribution. In the case of binomial data, the mean value of the Jeffreys prior is 0.5 (since the distribution is symmetrical between zero and one), which has a tendency to skew the *posterior* uncertainty distribution toward the middle (Atwood, 1996). For the constrained noninformative distribution defined by Atwood, the posterior uncertainty mean value tends to follow the user-specified mean value, as we would expect in the case of expert judgement.

4.1.2.4 Uncertainty in the Safety Assessment

The portion of the decision framework that generally has the most rigorous treatment of uncertainty is in the safety assessment. Most PRAs, and associated software analysis tools, have long been able to evaluate epistemic uncertainties related to the model parameters. Unfortunately, other sources of epistemic uncertainty such as modeling issues have not been investigated to the level of detail as that for parameters. These other epistemic areas should be considered in future research, but they were not the focus of our work. Further, we will not discuss in detail uncertainty related to PRA since the literature is replete with such discussion (McCormick, 1981; Henley and Kumamoto, 1981; Modarres, 1993)

As already mentioned, PRAs utilize two types of models, deterministic and aleatory. Within these models we have epistemic uncertainties. In the safety model, epistemic uncertainty deals with our state-of-knowledge about the components, safety system, and accident sequences. Model uncertainty falls under epistemic uncertainty but in the current decision advisor framework, we only consider parametric uncertainties. Within PRA, these parametric uncertainties typically fall into one of two categories, those related to a binomial process and those related to a Poisson process. From these process we can derive probabilistic distributions on the parameters of interest. For example, a beta distribution is frequently used when the failure is based on demand-type events (a binomial process). Recall that the binomial model for failures is defined as:

$$P(r) = \frac{n!}{r!(n-r)!} p^r (1-p)^{n-r} \qquad (38)$$

where r is the number of failures, n is the number of trials (or demands), and p is the probability of a failure per trial.

If the failure is based upon time-to-operate events (e.g., Poisson), then the gamma distribution is frequently used for the epistemic uncertainty. For example, the failure of a diesel generator to function for eight hours may utilize the gamma distribution to characterize its epistemic uncertainty. Recall that the Poisson distribution for failures is defined as:

$$P(r) = \frac{e^{-\lambda t}(\lambda t)^r}{r!} \tag{39}$$

where r is the number of failures, λ is the rate of failures (per unit time), and t is the duration of interest (eight hours in our diesel generator example).

These two probabilistic models, the binomial and Poisson, provide the backbone to the uncertainty analysis in most PRA and reliability modeling activities. These, coupled with deterministic models such as fault trees and events trees, provide the general framework of PRA.

4.1.2.5 Uncertainty in Economics

As discussed previously, the information for the economics module we have developed as part of the incident management prototype utilized (primarily) the NRC-developed *Regulatory Analysis Technical Evaluation Handbook*, NUREG/BR-0184; *Economic Risks of Nuclear Power Reactor Accidents*, NUREG/CR-3673; and *Generic Cost Estimates*, NUREG/CR-4627 (U.S. NRC, 1997; Burke, Aldrich, and Rasmussen, 1984; Sciacca, 1989, respectively). It is unfortunate though that the uncertainty treatment is superficial in both NUREG/BR-0184 and NUREG/CR-3673 and is nonexistant in NUREG/CR-4627. The data values in both reports are generally identified as "mean values" and will be used as such in the prototype development. Estimates of the epistemic uncertainty of these values are not provided, with one exception. In the case of severe accident related worker dose, a low and high value are provided. But, the remaining economic parameters (replacement power costs, worker salaries, discount rates, etc.) do not have a specified range of values.

If we turn to general, non-nuclear cost indicators, we can determine an approximate variation that one may expect to see for cost estimates. For example, evaluating the 1992-2001 consumer price indexes (U.S. Government, 2002), we can calculate that the aggregated price index for commodities and services varied, from year-to-year, by about 20% (as measured by one standard deviation). If we focus just on either commodities or services, costs associated with these categories were found to vary by about 30%. More volatile categories, such as energy prices, were found to vary by about 240%. Consequently, it is not unreasonable to assume that economic estimates will vary 50% about the point estimate value. We utilized this variation, with the assumption of a uniform distribution between the lower and upper points, for all cost parameters.

4.1.3 Example Epistemic Uncertainty Analysis

It is key to keep in mind that all of the discussion relevant to uncertainty must ultimately focus on the intended application of the decision advisor prototype. In short, the goal of the advisor system is to suggest, from a list of potential decision alternatives, which one is preferred. Here, preference directly implies use of expected disutility via Equation 37. But, it is not enough to simply perform Monte Carlo analysis on individual decision options, plot the percentiles and statistical moments, and from those make a decision. Instead, one must carefully consider the fact that the decision alternatives are not independent, and as such, can not be treated as if they were independent. Thus, what we need to calculate is the predictive distribution for the decision alternative rankings (i.e., which decision is "first," which one is "second," etc., and their associated probabilities).

As a numerical example of the uncertainty on decision ranking, we ran a case study for a leaking steam generator tube. We focused on four potential decision alternatives, Decision I was to stay at power and provided normal makeup water, Decision II was to shut the plant down and plug the leaking tube, Decision III was to reduce power to mitigate (some what) the leak, and Decision IV was to isolate (if possible) the leaking steam generator. For this analysis, we utilized an influence diagram and associated decision tree and put representative epistemic uncertainties on all parameters. Further, to help focus the analysis discussion, we utilized just cost consequences in this example calculation. The results of this analysis are shown in Figure 46. In this figure, we see that when considering the parameter uncertainties, that Decision I (use makeup and stay at power) is preferred about 99% of the time. Decision II (go to shutdown) is the preferential decision only 1% of the time. Also, note that Decision III (reduce power) and Decision IV (isolate SG) are never preferred.

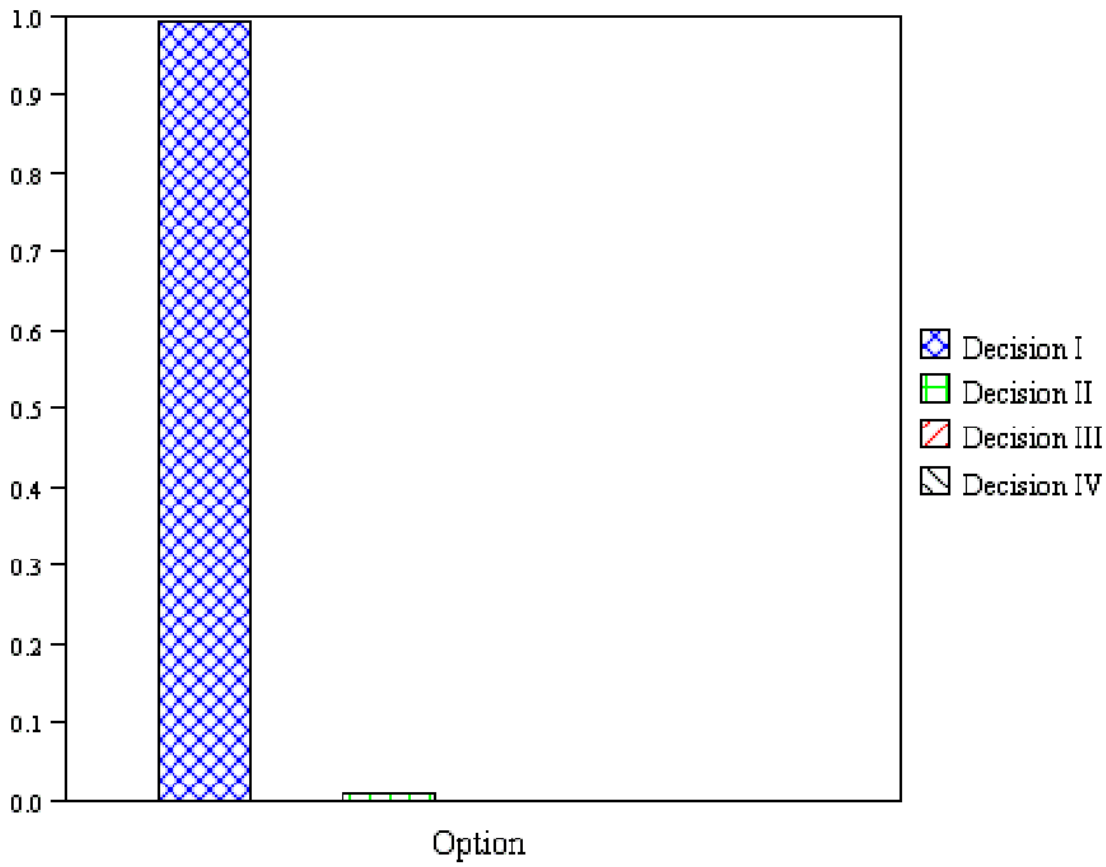

Figure 46. Monte Carlo sampling of the epistemic distributions for the leaking steam generator example parameters.

We further looked at the epistemic uncertainty with the decision making model by calculating the uncertainty on each decision alternative individually from each other. In other words, we performed the same calculation as that shown in Figure 46, but this time we stored that expected value of the *cost* rather than the decision *ranking*. The results of this uncertainty evaluation on these decision nodes are shown in Figure 47.

Looking at Figure 47, we note two items of interest. First, each decision alternative outcome has significant uncertainty. For example, focusing on Decision I, the 90% probability interval ranges from 500,000 to 20,000,000 euros. While Decision I does not have the largest uncertainty, it is representative of the uncertainty range on expected value that can be seen on a decision alternative. Recall that the outcome out of a decision model consists of a combination of chance variables (e.g., probability of a leaking tube leading to a severe accident) and cost variables (e.g., plant shutdown costs, repair costs). It is this combination of uncertain probabilities and uncertain outcomes that leads to variation in the expected value calculation.

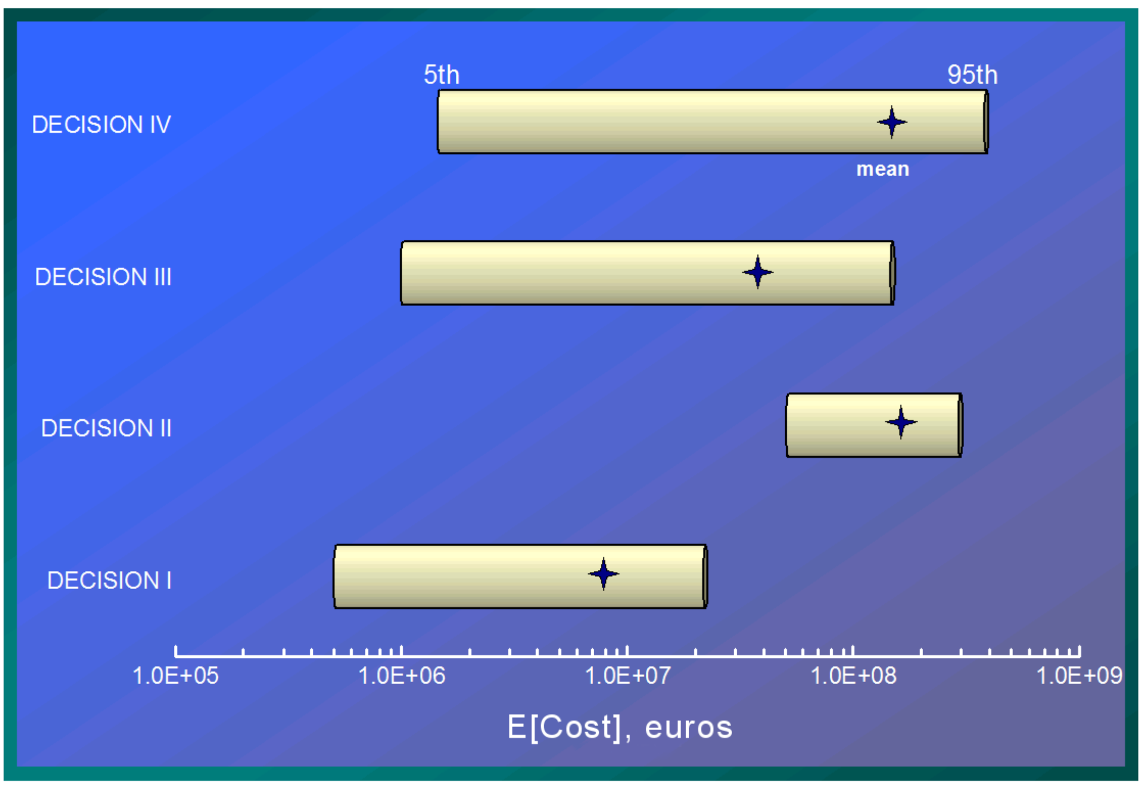

Figure 47. Monte Carlo sampling of the epistemic distributions for the each decision alternative in the leaking steam generator tube example.

The second, but most important, item to note from Figure 47 is that for each of the decisions, its 90% uncertainty range (i.e., the span of the bar shown from the 5% to the 95%) overlaps many of the other decision distributions. This concept of one uncertainty distribution (or range) overlapping another is known as stochastic dominance, where the degree of the overlap dictates the "dominance" of one decision over another. What is not shown in the figure is the degree of overlap since the decision alternatives are not probabilistically independent. One must be very careful not to be mislead by the "overlapping bars" shown in Figure 47 – the fact that two bars overlap to a large degree (e.g., Decision III and Decision IV) does *not* imply that that are approximately equal with respect to expected value. To explore this potential pitfall, let us focus solely on Decision I and Decision IV since they overlap to a large degree, but it was found (from Figure 46) that Decision I is *always* preferred over Decision IV, where preference is based upon choosing decision options that reduces cost.

We see from Figure 47 that the expected value for Decision I was lower than the mean value for Decision IV. Based upon this single metric, Decision I was the preferred decision. But, if we focus on the epistemic uncertainty of these two decisions, we see that there is significant overlap. In fact, we see that the 90% uncertainty range for Decision IV overlaps almost 80% of the uncertainty range for Decision I. One may be mislead into thinking then that since the two uncertainty regions overlap to a large degree that there is some fraction of the time where Decision IV would be preferred over Decision I. This line of reasoning is flawed due to the fact that the parameters for the calculation of expected value of Decision I and Decision IV are not independent. In fact, the parameters are, to a large degree in this example case, identical between the two decisions.

If we look at what makes up the two decisions (based upon a simple decision tree, which we do not show here), we see that each is a linear combination of parameters (consisting of chance variables and associated cost outcomes). The equations used by the decision tree are shown below:

Decision I = $C \cdot S \cdot \text{core\_damage\_cost} + C \cdot /S \cdot P \cdot \text{core\_damage\_cost} +$
$C \cdot /S \cdot /P \cdot (\text{leakage\_rate} \cdot 75000 \cdot \text{time\_to\_scheduled\_shutdown}) +$
$/C \cdot S \cdot \text{core\_damage\_cost} + /C \cdot /S \cdot P \cdot \text{core\_damage\_cost} +$
$/C \cdot /S \cdot /P \cdot (\text{leakage\_rate} \cdot 75000 \cdot \text{time\_to\_scheduled\_shutdown})$

Decision IV = $S \cdot \text{core\_damage\_cost} + /S \cdot P \cdot \text{core\_damage\_cost} +$
$/S \cdot /P \cdot (\text{shutdown\_cost} \cdot 0.33 \cdot \text{time\_to\_scheduled\_shutdown})$

where

C	=	Probability that makeup water fails
/C	=	Probability that makeup water does not fail
S	=	Probability that secondary cooling fails
/S	=	Probability that secondary cooling does not fail
P	=	Probability that primary cooling fails
/P	=	Probability that primary cooling does not fail
core_damage_cost	=	Cost if core is damaged (euro)
leakage_rate	=	Leak rate (liters/hr)
shutdown_cost	=	Replacement power cost per day (euro)
time_to_repair	=	Days to repair a leaking tube (days)
time_to_scheduled_shutdown	=	Time until the next scheduled shutdown (days)

As we see from the two equations above, many of the same parameters are found within the two decision outcomes. Consequently, any uncertainty analysis that compares the two decisions must consider the fact that these two decisions are correlated. Further, this correlation between decision alternatives will be the present in most incident management situations since the potential decision options available at the plant are limited in number and are somewhat similar in nature. While it is true that each decision alternative has substantial uncertainty in an absolute sense, one should focus on all the decisions as a whole during the process of decision making. In other words, the integration shown in Equation 37 must be performed outside the summation of the *i'th* performance measure and must encompass all potential decision alternatives. Note also that this problem is experienced in other applications, such as PRA importance measure determination, when sensitivity calculations are performed and subsequently compared.

Lastly, when we evaluate Decision I and Decision IV together, including the epistemic uncertainty, we find that Decision I is preferred over Decision IV 100% of the time (it shows complete stochastic dominance over Decision IV). One may wonder how this outcome arises when there is significant overlap between the two uncertainty distributions. The answer to this query is that one need to remember that the decisions are correlated, so when one of the parameters in Decision I has a high value (from the Monte Carlo sampling), the parameter has that same high value for the Decision IV case. Thus, given the formulation of the decision tree for Case 2, it turns out that Decision IV is never preferred over Decision I (for the nominal case). Note though that while Decision IV is never preferred over Decision I, this is not the case for *Decision II*. The equation for Decision II (not shown) has parameters that are independent of those in Decision I. Consequently, there is a chance that Decision II could be preferred over Decision I, even though the expected value for Decision II is much larger than the expected value for Decision I. But, after running the uncertainty analysis for both decisions, we found that Decision I is preferred over Decision II about 99% of the time.

4.2 Deterministic Analysis within the Decision Analysis Framework

The overall formal decision process is rooted in deterministic models such as the PI, value tree, disutilities, and fault trees from the PRA. We will not discuss the details of these models since the literature addresses much of the required background for value trees and disutility (Clemen, 1996; Keeney and Raiffa, 1993; von Neumann and Morgenstern, 1944) and fault trees (McCormick, 1981; Modarres, 1993). But, we will discuss attributes of the deterministic modeling as it relates to the PI, specifically discussing PI impacts for both plant operation and economics.

4.2.1 Deterministic Analysis in Plant Operation

The decision advisor prototype utilizes changes in plant status as a driver for a variety of impacts such as cost and external attention. Embedded in the analysis solution for the decision model are heuristics that map upsets such as transients into observable PI outcomes. To categorize these upsets, we utilize the initiator designations from our decision maker's PRA. In Table 17, we list these categories along with the initiator impact on our performance measures of cost and external attention. The values contained in this table were elicited from our decision maker and are used as the default values for the associated decision model parameters in the decision advisor prototype.

Table 17. Operational impacts to the cost and external attention performance measures.

Initiator	Plant down time (days)	Special repair costs (euro)	External attention outcome
Loss of coolant break	90	20,000,000	Long shutdown
Secondary-side pipe break	90	5,000,000	Long shutdown
Steam generator tube rupture	90	5,000,000	Ultimatum
Main steam line break	90	5,000,000	Long shutdown
Anticipated transient without scram	30	2,000,000	ultimatum
Loss of heat sink	30	2,000,000	ultimatum
Loss of electric power	8	none	Inspection (nominal)
			Ultimatum (if common failures of the diesel generators)
Secondary-side transient	1	none	Report
Primary-side transient	1	none	Report

We will see later that the initiators are also used to determine plant state probabilities via aleatory analysis. Consequently, these initiating events end up being used in two ways, first to determine the probability of getting to a decision-specific state and then second to determine the outcome of that state.

4.2.2 Deterministic Analysis in Economics

The NRC has expended considerable effort collecting information related to costs of certain activities at nuclear power plants. For our incident management prototype, we utilize much of this work (U.S. NRC, 1997; Burke, Aldrich, and Rasmussen, 1984; Sciacca, 1989; Claiborne et al, 1989; Lopez and Sciacca, 1990) in order to determine costs associated with decision alternatives and outcomes. For the economics modeling, the key tenants to the analysis are to (1) adequately identify primary cost factors and (2) discount costs as necessary.

Expenses specific to decision alternatives should consider a number of factors. For example, the total cost of a decision outcome is the sum of many different types of expenditures, possibly including: (Lopez and Sciacca, 1990)

1. Equipment and material costs.
2. Labor costs associated with installation and/or removal.
3. Costs associated with engineering and quality control and quality assurance.
4. Personnel staffing levels due to radiation exposure (i.e., spreading the dose amongst personnel).
5. Costs to defuel, drain, and restore the reactor.
6. Replacement power costs.

We will briefly address the important categories below.

Labor, Equipment, and Material Costs

The NRC provided the data which may serve as the basis for the equipment costs, material costs, and labor estimates (Claiborne et al, 1989). The NRC data incorporates "as-built" cost information for U.S. nuclear plant activities. For operating plants, there are a number of workplace features which may impact the level of productivity and thus increase the number of labor hours required to accomplished a task. But, these characteristics, would be specific to incidents that are being analyzed and would need to be indicated by the user of the incident management prototype. Also, the labor costs associated with activities include overhead charges, administrative support, rent, insurance, etc., and should be factored into the overall labor rate for the analysis.

Burke, Aldrich, and Rasmussen looked at repair costs for U.S. nuclear power plant incidents and determined an approximate cost of 1,900 euro per hour represents actual operating experience (1984), as indicated in Figure 48. We utilized this value as the nominal labor, equipment, and materials cost in the decision advisor prototype. Special equipment or staffing circumstances deemed beyond this average cost would be added to the average cost estimated by the prototype.

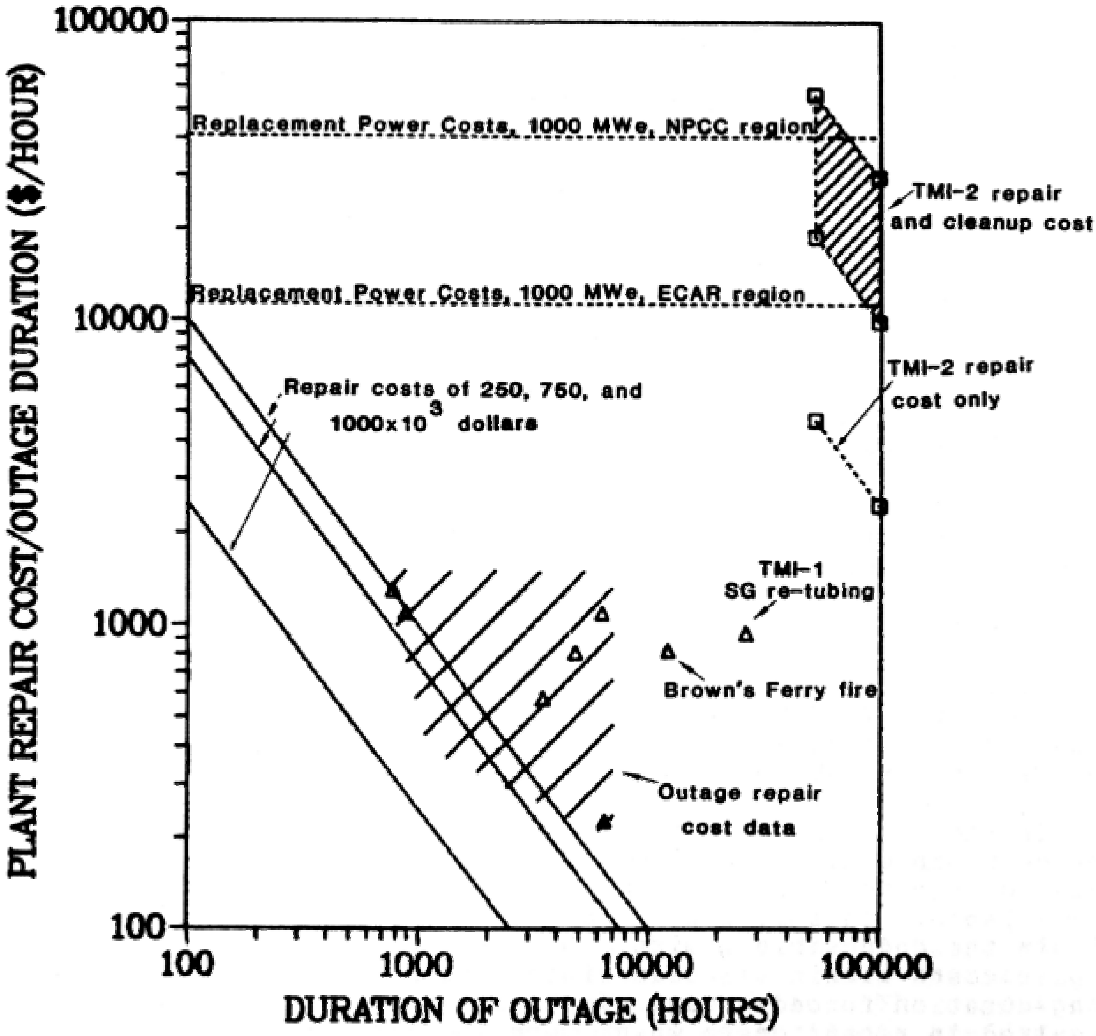

Figure 48. Example of nuclear power plant repair costs (from Burke, Aldrich, and Rasmussen, 1984) in 1984 U.S. dollars.

Radiation Exposure Estimation

Worker radiation exposure estimates will be dependent on the type of incident and the decision alternative considered. The collective radiation exposure associated with a decision should be estimated by multiplying the total labor hours necessary to perform the activity (in the radiation area) and the work area dose rate associated with the particular task. The work area dose rate will range from very low doses (negligible) to very high-dose areas (approximately 0.1 Sv/hour). We utilize the dose rate as a multiplicative factor on the total labor/material cost. If the dose rate is low, then the cost multiplier is one. A high dose rate is defined as a rate larger that 0.001 Sv/hr (U.S. Government, 2001). For rates between 0.001 Sv/hr and 0.1 Sv/hour, we assume that the cost multiplier is two.

Cost to Defuel, Drain, and Restore the Reactor

If a nuclear reactor core is left in place during shutdown operations, high radiation levels may be realized at certain locations of the containment area. Therefore, worker activities within these areas may be limited and would need to be factored into the potential decision alternatives. Complications or safety issues will lengthen the duration of the outage, where the cost of the outage can be estimated by knowing labor, material, repair, and replacement power costs. These impact must be specified by the user of the incident advisor prototype.

Replacement Power Costs

Replacement power costs for the potential plant modifications are one of the important factors in the overall methodology. Estimates for this category of cost were developed based on information provided to us by our decision maker. A best estimate of 333,000 euro per 24 hour period was indicated as an appropriate value. Note though that this cost value does depend on the specific time of year, regional power requirements, and duration of the outage, and as such, may vary from one incident to the next. To account for possible variations, the incident advisor prototype allows the user to modify base values with incident-specific information.

One interesting aspect of the decision process relates to the operation of a fleet of nuclear power plants. If our decision analysis methodology is used to determine a change that will affect multiple plants, one must ensure that the disutility analysis is performed correctly. For example, if an implementation cost associated with a decision is one million euro, but this decision is going to be applied to 20 plants, the disutility applicable to this activity is u(20 $\times$ 1,000,000 euro), *not* 20 $\times$ u(1,000,000 euro).

4.3 Aleatory Analysis within the Decision Analysis Framework

Much of the analysis within the decision framework takes place within the deterministic models discussed in the previous section. Nonetheless, analysis of the aleatory portion for the incident advisor is a critical part due to the fact the aleatory analysis determines the chance nodes of the decision model. This analysis will dictate the likelihood of reaching outcomes such as a particular plant state (e.g., transient, loss of power, core damage) or impacts to workers (e.g., injuries, fatalities). In Figure 49, will illustrate a selection of these chance nodes that are germane to the generic influence diagram structure of the decision advisor.

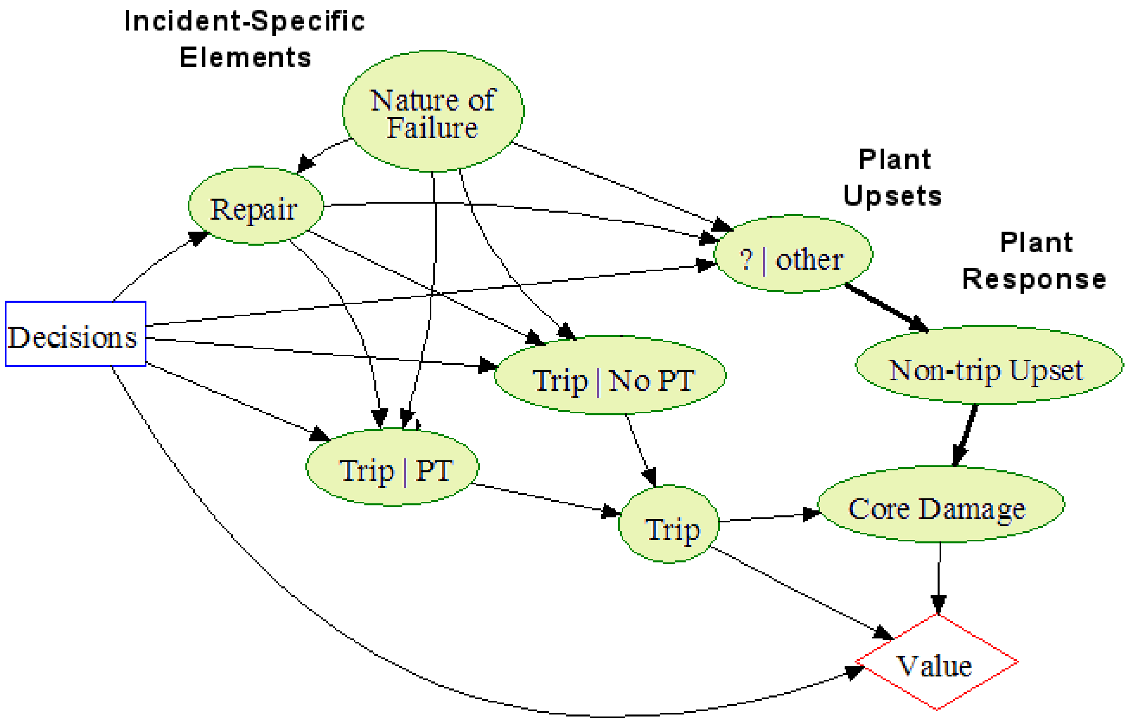

Figure 49. Example generic influence diagram with the chance nodes highlighted.

4.3.1 Aleatory Analysis in Plant Operation

Nuclear power plants are designed to respond to a variety of upset conditions. These initiating events have an estimated frequency of occurrence, from the PRA, and a corresponding conditional core damage probability (CCDP), also from the PRA. To calculate the CCDP for the occurrence of an initiating event, one must modify the PRA to account for the initiator that happened and the ones that did not. Then, any complications that are part of the incident, for example a failed component that is not restored, must be indicated in the PRA. In general, three items must be considered before the PRA can be used to estimate a CCDP.

1. Finding the appropriate basic events in the PRA in order to change their failure probability (for failed components) or initiator frequency to reflect the specifics of the incident.

2. Assessing the probability that an initiator or affected component could be recovered (i.e., returned to an operational state). For example, if a small break loss-of-coolant accident were to occur, the chance of isolating the leak may be large and should be factored into the PRA calculation.

3. Determining other impacts on the PRA model such as changes to common-cause failure probabilities (due to a failure) or component alignment/operation issues.

To model the aleatory core damage impact of an initiator, we first setting the initiating events to either probabilities of one (if it happened) or zero (for the remaining ones). Then, one would normally consider the potential of recovery from the initiator. But, for the decision advisor prototype, we are not going to credit initiator recoveries since the plant upset will occur and the impact to the PI performance measures may be felt, regardless of the initiator degree. Other component-type recoveries embedded in the PRA will be left as is since the operators may be required to restore failed components following an initiating event. If the incident itself is one such that a component is failed, then its related impacts to the fault and event tree PRA models are included during the CCDP calculation. After the PRA model is adjusted to reflect both the initiating event and the component failures (if any), then the CCDP determination may proceed. Smith (1998) and François (1998) provide additional information and insights for CCDP calculations of nuclear power plant risk models.

From the decision maker's PRA, we have determined nominal initiator frequencies and CCDPs (for the full-power mode, state A). These values, shown in Table 18, are embedded in the knowledge base utilized by the decision advisor prototype since they represent the outcomes identified in our generic chance nodes.

Table 18. Nominal initiator frequencies and impacts to the core damage probability (assuming no other component failures).

Initiator	Nominal frequency (per hour)	Nominal CCDP (plant state A)
Loss of coolant break	2.9E-3	6.1E-4
Secondary-side pipe break	1.3E-3	7.1E-4
Steam generator tube rupture	1.6E-2	3.3E-5
Main steam line break	4.2E-3	2.1E-4
Anticipated transient without scram	7.4E-1	1.9E-7
Loss of heat sink	1.3E-4	3.0E-3
Loss of electric power	6.7E-1	1.1E-6
Secondary-side transient	7.0E-1	7.5E-8
Primary-side transient	3.9E-2	7.1E-6

We have chosen to model plant upset conditions as leading to impacts in the PI performance measures such as core damage, cost, and external attention. The other two remaining performance measures, industrial accidents and radiological dose, are impacted by worker operations. These two measures represent a health burden on workers, and as such, are modeled via a separate (from the PRA) model. But, the work impact approach is similar in that we determine the rate of impacts (like initiators are impacts to the physical operation of the plant) for items such as injuries and deaths via an aleatory process.

For the determination of worker accident rates, we utilized data collected for the U.S. Department of Energy Computerized Accident/Injury Reporting System (CAIRS). CAIRS is a centralized point used to gather reports of injuries and other accidents that happen during Department of Energy operations. Worker accident data from 1995 to 2000 have been collected, analyzed, and processed for the ESHRAP risk analysis software developed by the Idaho National Engineering and Environmental Laboratory (Eide and Wierman, 2002). The CAIRS data set provides both the number of worker hours and the number of incidents, thereby allowing for case rates to be estimated. For the incident advisor prototype, we utilized the categories of rates of worker incidents, the probability of minor injury (given a worker incident), the probability of severe injury (given a worker incident), and the probability of death (given a worker incident). Note that we defined minor injuries as those that resulted in a work restriction but no lost days away from work. The CAIRS data results are listed in Table 19.

Table 19. U.S. Department of Energy CAIRS worker accident data.

Case	Nominal frequency (per hour)	Nominal probability
Worker incident	1.6E-4	
Negligible impact		5.34E-1
Minor injury		2.57E-1
Severe injury		2.09E-1
Fatality		5.0E-4

4.3.2 Aleatory Analysis in the Safety Assessment

The plant safety assessment, as defined by the PRA, addressed component failures, human actions, and system response through a variety of models. Earlier, we identified where in the decision model (Figure 49) many of these impacts are felt. The aleatory analysis utilized from the PRA includes the determination of failure probabilities for degraded systems and operator actions, both prior to an upset condition and following an upset condition. For information into degraded system modeling using the PRA, we refer the reader to Smith (1998).

One of the decision alternatives that may be frequently encountered is the potential of repairing or replacing an inoperable component without shutting the plant down. Thus, in the knowledge base behind the incident advisor, we have defined a "decision node" representing this decision alternative. But, in addition to this node, we have identified two other nodes that represent the repair process. The first of these two nodes contains an aleatory model that determines the probability that the repair process causes an upset condition (e.g., an operator inadvertently trips the plant by shorting a circuit). The second of these two nodes contains an aleatory model that determines the probability that the repair will be successful or not. Both of these two nodes require determination of human performance.

One way to obtain an estimate for the human failure probability (such as a test-caused failure) is to use an existing database or a methodology for estimating human error rates. For example, the Technique for Human Error Rate Prediction could classify incident-related events (Swain and Guttmann, 1983). Or, other methods, such as the Accident Sequence Evaluation Program (Swain, 1987) or the checklist method for the human reliability portion of the U.S. NRC Standardized Plant Analysis Risk models (Smith et al., 2002) could be used.

The Accident Sequence Evaluation Program does provide incident-type generic human error rates. For example, the probability of human failure for preventive maintenance or tests, assuming no recovery factors, is given as 0.03. This value is considered to be a rough estimate for operators not effectively performing a test or maintenance.

Additionally, the NUCLARR database could be used to obtain a probability of test-caused failures (Gertman et al., 1989). From this database, the test-caused probability of failure was found to be 0.002, which represents an aggregate of a variety of failures. In nuclear power plants, the use of procedures by plant personnel minimizes the chance of having a test-caused failure. But, for complicated systems or complicated testing procedures, the probability of a test failure may be larger than the 0.002 average value. Conversely, for simple systems or uncomplicated testing procedures, the probability of a test-caused failure may be much lower. In the incident advisor prototype, we allow the user to specify an incident specific value for the the human errors.

For complex testing arrangements, where components and system states change during the procedure, a static failure model may only provide a rough estimate of the failure probability. For more detailed analyses, on could utilize Markovian state change analysis. This technique requires that all potential states (operating, under repair, being tested, etc.) of a particular system be known. As such, the number of potential states may become large as the number of components in the system increases and the number of possible states for each component increase. It is this potential for a large number of system states which results in intractable analysis for many Markovian analyses. While past research has attempted to overcome the limitations of Markov analysis for systems such as nuclear power plants (Papazoglou and Gyftopoulos, 1977), Markovian analysis is not routinely employed. But, if the delineated states are few, Markov techniques may be useful to represent dynamic situations.

Related to Markovian analysis is the technique of event simulation. While not typically used as part of formal decision analysis, simulation may be a powerful analysis tool since, for any one decision alternative, the state of the plant may take a variety of different paths depending on the aleatory outcomes. For example, if one decides to continue operating the plant with a degraded safety system, the plant may (a) function properly for a lengthy period; (b) undergo an upset condition and then stay shut down due to the upset; or (c) undergo an upset and quickly return to full operation. A static model, such as a decision tree, may only approximate these plant states and their associated impacts on the PI decision criteria. Simulation allows us to bring a higher fidelity to the analysis process, but at the cost of longer analysis times. Others have discussed (for static models) the fact that lengthy analysis times poses problems for real-world use of decision analysis techniques (Call and Miller, 1990).[†] Nonetheless, the idea of wrapping the entire decision process, from the time of the decision until the final outcome at a future point in time, in a simulation framework deserves attention – toward that end we devote a section to explore associated issues and techniques.

4.4 Process Simulation

One of the novel ideas expressed in this document is our approach to solving the decision model. While static solution techniques such as the "roll-back method" (Clemen, 1996) of decision trees may work in many situations, there may be other situations that require a more robust analysis regime. Thus, we proposed to decompose the entire decision process in a simulation framework. While previous researchers have utilized what they called "simulation" to sample the epistemic distributions of decision attributes, we note that a more suitable nomenclature for this analysis is "Monte Carlo uncertainty propagation." When we use the term "simulation" in this thesis, we intend its meaning to be the classical definition, namely that of stochastic simulation to produce results from a model that has stochastic elements (Ripley, 1987), where here stochastic has the same meaning as aleatory. Of course our simulation will have unknown parameters – or

[†] Thanks to modern computers, our test PC (1.5 GHz Pentium 4) is able to provide approximately one million random numbers every three seconds. This calculation speed indicates that the issue of "fast enough" is becoming less of a problem in the decision analysis domain.

epistemic uncertainties – which we propagate through the simulation via Monte Carlo methods. But the simulation itself deals with the “randomness” one faces following a decision, randomness that tells us that we will not know exactly how the “play” ends even though we have a good idea of the actors and the story.

4.4.1 Why do we need simulation?

As part of the work for this document, we evaluated several incident case studies. We used these analyses as a test-bed for ideas to help guide the research. One of these case studies evaluated decision making following a failure of a component that senses water pressure in the primary coolant loop of a nuclear reactor. This component, a pressure transducer, is only one of four such components. Consequently, we can operate the plant while one of these components is failed. But, our *initial* decision model for the pressure transducer did not include the potential for repair. An important question then arises; "how realistic is the no-repair assumption with respect to the decision model?"

If the decision maker chooses the decision alternative of continued operation of the power plant (with the pressure transducer failed), then he or she runs the risk of facing an inadvertent plant trip (from one of the three other transducers). While the possibility of a single trip is captured in the pressure transducer model, there still exists a modeling limitation since it is possible to trip, repair the component, and then continue operation (we presume that the plant would not restart until the failed component was repaired). But, a second decision alternative for this model could be to shutdown and fix the problem immediately. Thus, we have, at first glance, different time duration for each alternative. Note though that we should have a consistent duration for the decision alternatives, otherwise we are comparing two options that are not really comparable. To illustrate this point, let us look at a portion of the initial decision model related specifically to tripping the plant via the pressure transducers. We will focus just on the chance node representing the trip (noted as "Spurious Shutdown" in the chance node) and the outcome (noted as "Loss" in the value node) in Figure 50.

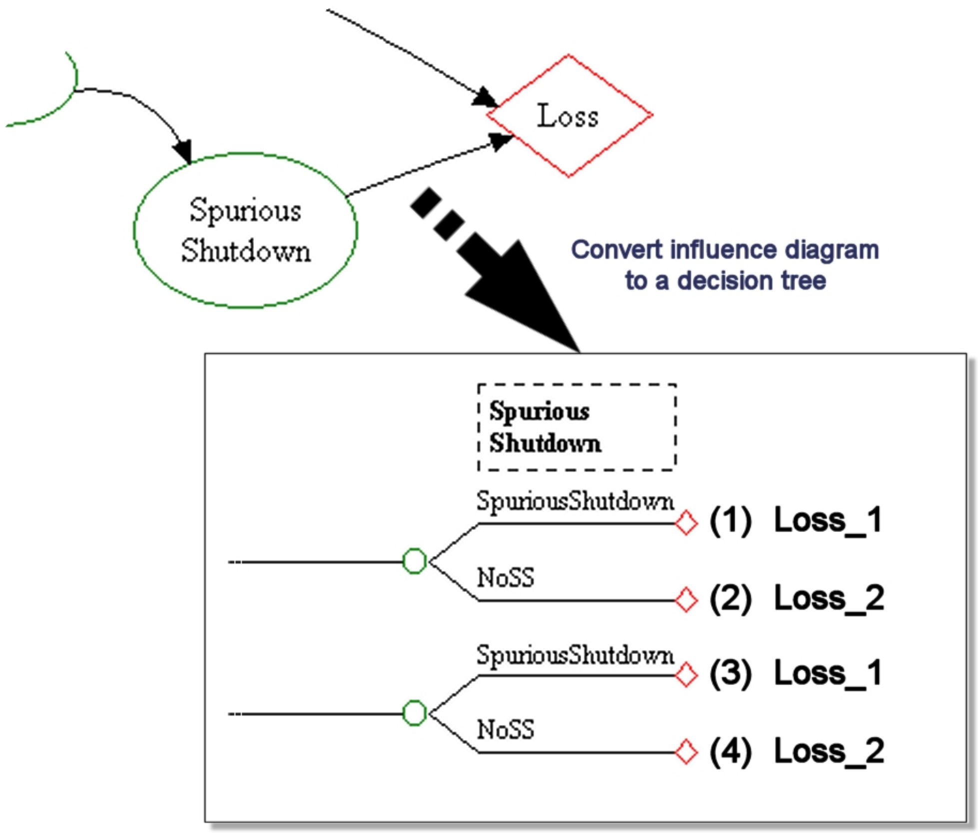

Figure 50. Example nodes from the initial pressure transducer case study decision model.

As noted in Figure 50, the outcome of a situation where the plant undergoes (one or more) spurious shutdowns is given by the variable "Loss_1." The other outcome, "Loss_2," represents situations where the plant does not trip (with probability given by "NoSS") due to the pressure transducers and will not be discussed further. The outcome represented by Loss_1 represents the PI for the decision tree sequence of interest, where in this case, we are focusing on either sequence (1) or sequence (3), depending on the decision alternative. Recall that the PI is a multi-attribute metric which factors in the decision makers weights (on the *i'th* PI), the disutility for the outcome, and the probability of seeing that outcome. For example, sequence (1) may include impacts such as costs (due to the plant tripping inadvertently), increased external attention (perhaps by negative public following the plant trip), and an increase in the CCDP (since we are challenging safety systems after the plant trip).

Complicating the determination of the Loss_1 outcome is the fact that the decision outcomes may span a long period of time. Nominally, the time period of interest is the duration until the next scheduled plant outage, which could be as long as 18-months for current nuclear power plants. But, for some issues, it is possible that the time window of interest could be measured in years. Thus, we need to consider the possibility of a scenario like:

1. The plant trips inadvertently
2. The degraded pressure transducer(s) are repaired
3. The plant is restored to operation
4. The plant trips sometime after returning to operation

If the duration is sufficiently long, one may need to consider multiple plant trips. Contrast this scenario with our original pressure transducer model. The original model did not include the potential for repair following a trip, but we just indicated that repair following a plant trip is what would normally happen. Further, the original model allowed a maximum of one plant trip over the entire duration, regardless of the length of the duration. Clearly, this simplifying assumption only would be valid for those cases where it is unlikely to have the pressure transducer trip the plant. Others have noted that static models have constraints that limit their usefulness when compared against simulation techniques (Goel and Ren, 1999).

The outcome represented in the decision model must exhibit a one-to-one relationship with the supporting models. In this case for the pressure transducers, if the plant experiences more than one trip, the PI for the outcome node *must* reflect the multiple outages. If each outage following a plant trip results in several days of lost power generation, the cost impact could be quite large. Further, if a plant experiences more than one trip in a relatively short period of time, other impacts like external attention may be increased due to the sensitive nature of initiating events at nuclear power plants. The original static model does not reflect these time-dependent situations.

4.4.2 Simulating Dynamic Processes

Since the standard decision tree and fault tree model may not represent our decision process adequately in some cases, we can turn to more sophisticated analysis techniques. Fortunately, the nuclear power plant PRA research community has explored related analysis issues, namely on the topic of "dynamic PRA." Here, the term "dynamic" is intended to bring in the aspect of time and the evolution of plant processes as an integral part of the PRA model. For an overview of several dynamic methods related to PRA (circa early 1990), we direct the reader to Siu's review paper (1994). From the research literature, we were able to distill the major elements of the PRA dynamic methods into a few salient features, namely how the methods treats time and model space. For this body of work, we show a summarization of the major modeling methods in Figure 51.

A few words about the information in Figure 51 are necessary. First, all the reviewed dynamic methods proposed for PRA applications are a subset of general Markov theory (for example, see Howard, 1971). Consequently, we list the four major aspects of the Markovian theory behind these dynamics method: (1) renewal processes, (2) vector processes, (3) transient processes, and (4) semi-Markovian processes. With Markov theory as an underlying background, one can then attack the dynamic modeling problem by utilizing (1) discrete space (D-space) or continuous space (C-space); (2) discrete time or continuous time; and (3) Markovian or semi-Markovian state transitions. Note that very few proposed dynamic PRA models utilize both continuous time and space.

The simulation technique we are proposing as part of the decision analysis framework follows that of discrete event simulation. In general, discrete event simulation is a semi-

Markovian technique utilizing both discrete time and discrete space. In our case, time is discretized from the point of the decision until the next scheduled plant outage while space is discretized by way of the plant operational condition. The aleatory and deterministic models we described earlier allow us to determine the movement through the phase space of the simulation.

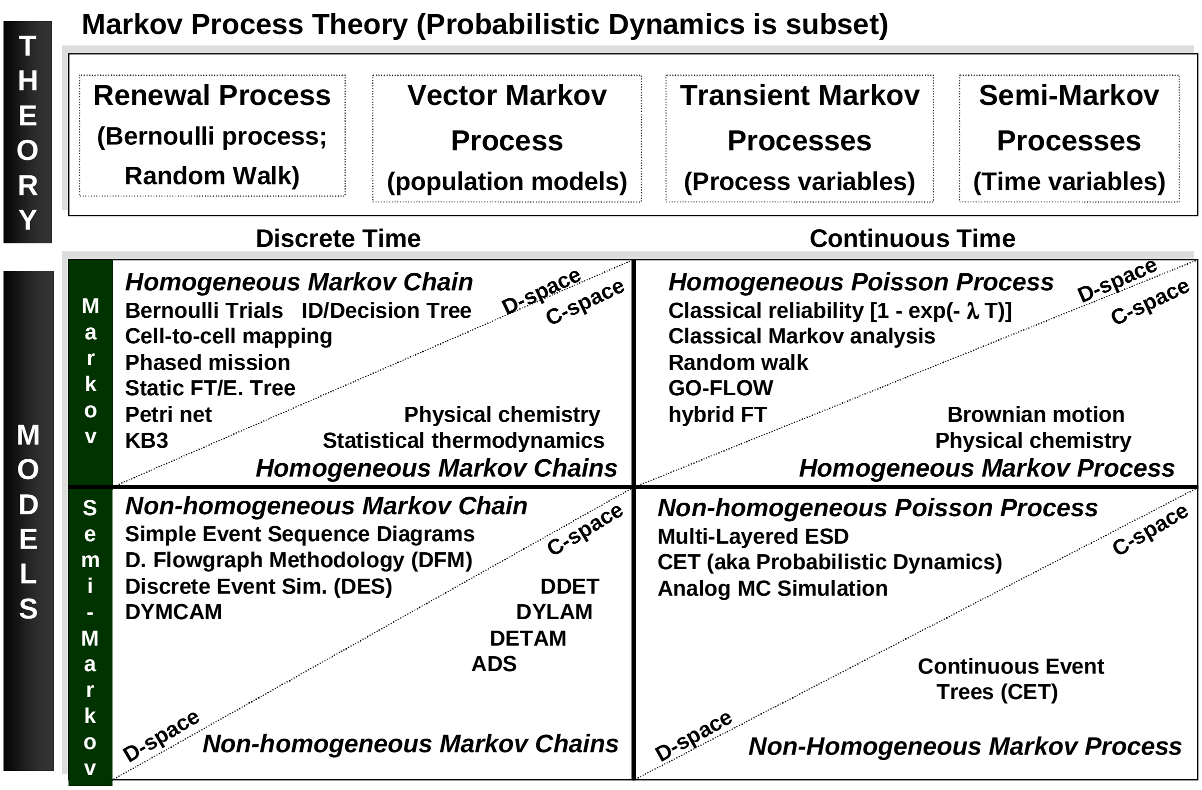

Legend: D-space = discrete space C-space = continuous space

Figure 51. Characterization of dynamic PRA analysis methods.

To better describe the use of simulation, we return to the pressure transducer example. A simulation was performed using the algorithm described in Appendix E for a 1-of-3 pressure transducer system, where a trip occurs if one (of the three) transducer fails. We utilized this algorithm to determine system performance where the component failure rate was varied from 1×10^{-5}/hr to 1×10^{-2}/hr, all for a fixed mission time of 2,400 hours. For this evaluation, we calculated an expected number of plant trips, but we could also determine a point-wise unreliability or the average unreliability. The simulation model was designed to include the potential of repair if the plant inadvertently trips. For this example, we assumed that the plant would be off-line for two days to repair the pressure transducer following the plant trip.

In order to check the pressure transducer simulation calculation, we need to determine an analytical solution to the 1-of-3 (failure criteria) system. The aleatory model that represents an individual transducer failure is given by the unreliability equation:

$$U(t)_{transducer} = 1 - e^{-\lambda T} \tag{40}$$

where λ is the failure rate and T is the mission time. We note that the component reliability for a pressure transducer is given by R(t) = 1 - U(t). For a system of *m* identical components, the *system* reliability is (Martz and Waller, 1991)

$$R(t)_{system} = \sum_{x=r}^{k} \binom{k}{x} R(t)^{x} [1 - R(t)]^{k-x} \tag{41}$$

where r is the success criteria (i.e., the number of components that must function for the system to function), k is the number of redundant components, and R(t) is the component reliability. In our case, k is three while r is three (for success), or:

$$R(t)_{system} = R(t)^{3} [1 - R(t)]^{0} = e^{-3\lambda T} \tag{42}$$

We can now determine the point-wise unreliability for the system:

$$U(t)_{system} = 1 - e^{-3\lambda T} \tag{43}$$

If λ is 2×10^{-5}/hr and T is 2400 hours, $U(t)_{PT\ system}$ equals 0.134. Note though, that the value represents a point-wise unreliability. This point-wise metric is the probability that one of the three pressure transducer failed (and was not repaired) at a time of 2400 hours. It does not indicate how much of the 2400 hours was spent in a "failed" state nor does it include the potential for repair. To determine the average unavailability, we must integrate the unavailability expression over the mission time, or:

$$E\left[U(t)_{\ system}\right] = \frac{\int_0^\tau U(t)_{\ system}\, dt}{\int_0^\tau dt} = \frac{\int_0^\tau 1 - e^{-3\lambda t} dt}{\int_0^\tau dt} = \frac{\int_0^{2400} 1 - e^{-3\times 2\times 10^{-5} t} dt}{2400} = 0.0687 \tag{44}$$

We then ran the system simulation using 20,000 iterations to compare against the analytic solutions derived above. We found, from the simulation, that the average unavailability was 7.1×10^{-2}, assuming a 1-of-3 (for failure) system, a failure rate of 2×10^{-5}/hr, a mission time of 2,400 hours, and no repair. The simulation results are in good agreement with that calculated via Equation 44 (it differs by about 3% from the exact calculation). Additional confirmatory calculations where we vary the failure rate may be found in Appendix E. Further, we used the simulation to determine the number of trips of the system. The simulation results for the expected number of trips are shown in Figure 52. Several observations related to these results must be noted.

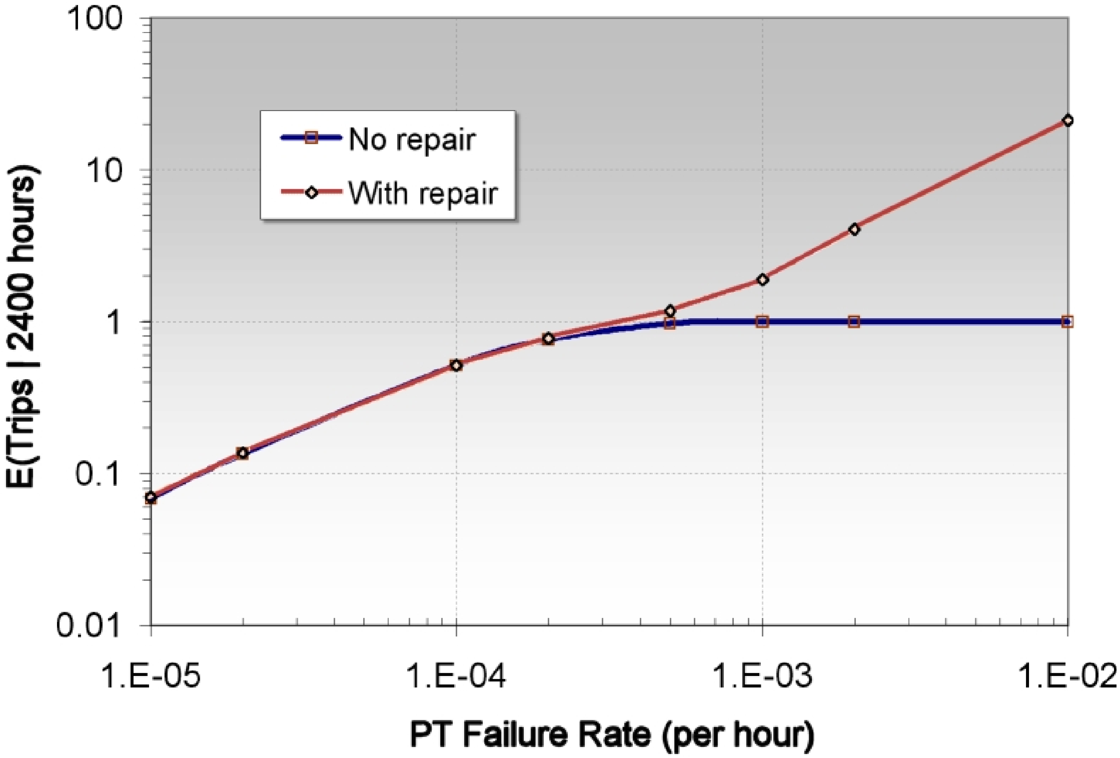

Figure 52. Simulation of system performance to illustrate the potential for multiple plant outages following component repairs.

First, for the case where repair is possible, an analytic solution to the problem would be difficult due to the fact that the reliability model considers two pressure transducer system states. The first state follows the initial failure of the pressure transducer (when the decision to continue operation is made) but is prior to an inadvertent plant trip. In other words, the first plant state represents a degraded 2-of-4 (for failure) system. The second state follows an inadvertent plant trip and returns to a non-degraded 2-of-4 system. Any analytic solution would have to consider these two states, and the transitions between them, in an integrated fashion. Of course, the simulation models both the degraded and non-degraded pressure transducer configurations as a direct part of the analysis, as we illustrate in Figure 53 for six iteration cases from zero to 2,400 hours (with a transducer failure rate of 1×10^{-3}/hr).

Second, for low failure rates (less than about 2×10^{-4}/hr), the two cases, with and without repair, have effectively the same expected number of trips. Further, this value (E[trips]) is equal to the probability of trip that was described by the original pressure transducer model we utilized. It should be realized that for any process where the expected number of upset conditions is low (E[upset condition] << 1.0), then static modeling will provide adequate probabilistic input into the decision model. But, if a static model is used, one must be careful not to violate the low probability assumption during an uncertainty or sensitivity analysis.

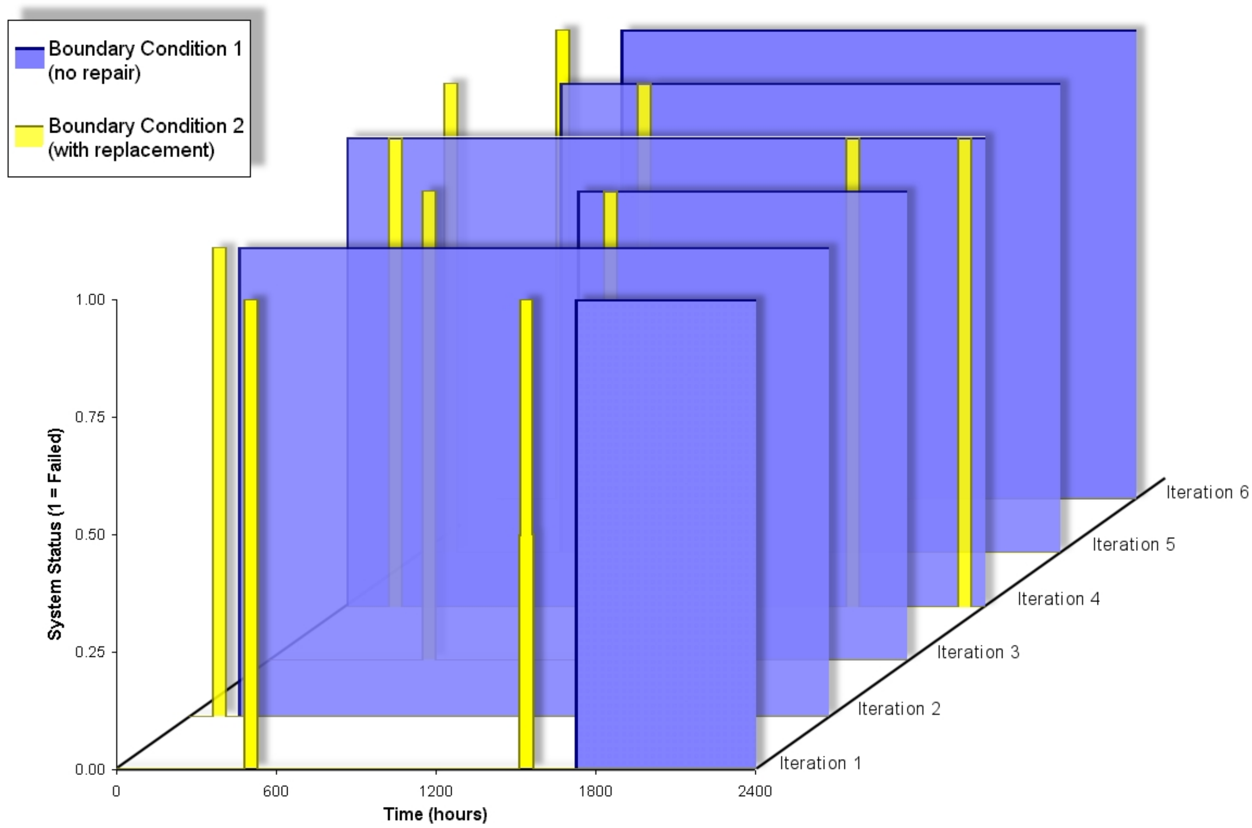

Figure 53. Six iteration results from the system simulation for the repair and no repair cases.

Third, the structure of the system, coupled with the component failure properties such as failure rates, determine the likelihood of spuriously tripping the plant. Note that as the failure rate of an individual pressure transducer increases, the expected number of trips increases. One should note from Figure 53 that, in some iterations, the plant tripped (due to the transducers) more than once since the failure rate was large (for illustration purposes). Adding redundancy to the transducer system would reduce the expected number of trips, especially for the case where one of the transducers has failed in a safe state.

Fourth, the epistemic uncertainty is not taken into consideration in Equation 44. Since the pressure transducer failure rate is not known, we would need to integrate over this variable also in order to correctly factor in its variability. But, in many cases, this calculation becomes analytically intractable. Conversely, to include the failure rate epistemic uncertainty in the simulation, we would only need to sample (via Monte Carlo) the failure rate a number of times (defined by the sample size) and repeat the simulation. Of course, this ease in calculation comes at the expense of calculation time.

A second simulation was performed to contrast the calculation precision between static models (such as fault trees) and dynamic models (our simulation). We evaluated a four pump system, where the pump failure rate was 1×10^{-4}/hr, the overall mission time was 8,760 hours, and we assumed that repairs were not possible. This case was run for such a lengthy mission time to illustrate the problems one may encounter when modeling aleatory behavior using static methods (such as fault trees) rather than simulation. For this analysis, we calculated metrics for both an individual pump train and the overall system in several configurations (i.e., a series 1-of-4, 2-of-4, 3-of-4, and parallel 4-of-4). For the pump train, we calculated the average unavailability and point-wise unavailability (at the end of the mission time), where these unavailabilities are represented by their classical definitions (Apostolakis and Chu, 1980). For the system configurations, we only calculated the system average unavailability over the mission time.

The results of that static and simulation analysis are shown in Figure 54. We provide three quantification techniques for the fault tree model: (1) the standard minimal cut set upper bound approximation and (2) an exact probability via the "inclusion-exclusion" technique. For each case, we plot the simulation results and then the exact results. Note that the exact results were determined by evaluating the analytic solution for system failure and integrating over the mission time to obtain the final results.

As we can see in the results, in some cases the fault tree results for a single pump train is too conservative. For example, the pump average unavailability is too large by a factor of almost two. If we look at the system results, we see that as the system unavailability decreases, the fault tree results become closer to the exact value. The simulation results match quite well with the exact calculation for all cases.

One should note that having an "exact" probability from a fault tree model would not necessarily yield the correct probabilistic answer. Since fault tree cut set solutions do not (generally) integrate over the time period of interest, it is difficult to determine an average unavailability using these models. Consequently, having advanced solution techniques, such as binary decision diagrams, to solve questionable models (e.g., static fault trees for dynamic scenarios) or to determine averages may provide little assistance.

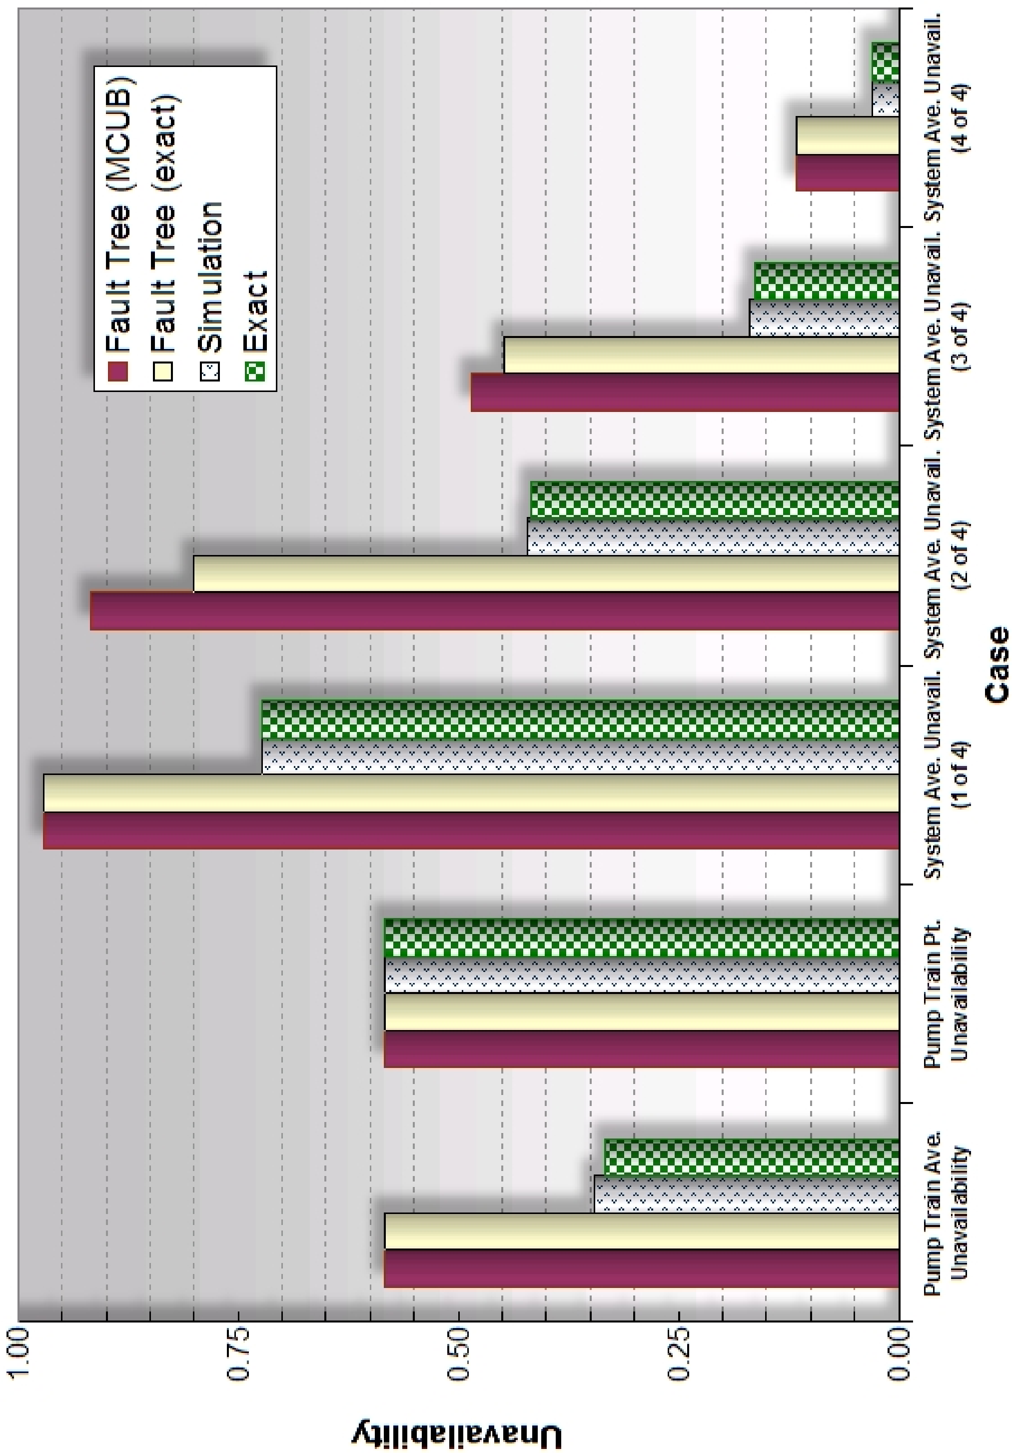

Figure 54. Results comparisons for fault tree analysis versus simulation.

4.4.3 Simulation of Decision Processes

The discrete event simulation approach to address the implementation and outcome of decisions relies upon traditional Monte Carlo methods to trace a sequence of (potential) events through time. As part of this simulation, impacts related to the value tree attributes must be considered. For example, if one iteration of the simulation leads to an extended plant outage, then the outcome will be felt though the cost and, possibly, through the external attention portions of the PI. The overall impact of this outage will be measured in terms of a vector of disutility attributes, including costs, worker impact, safety, and external attention. Thus, the disutility for this particular iteration will be higher than another iteration where no adverse situations were encountered. But, since these types of events are rare, the probability of seeing a long shutdown will discount the PI disutilities. The goal of the simulation is to run though as many potential outcomes as possible, determine their individual impacts, weight the impacts by their respective likelihood, and then determine an overall E[PI] for each decision alternative. Then, the decision alternative with the lowest E[PI] is considered to be the preferred option.

In order to implement the discrete event simulation, we could utilize one of two basic approaches. First, the simulation could question state transition potentials at each incremental time step. This method is referred to as simple or thinning event simulation since it is a straightforward (but inefficient) approach where potential state changes are allowed at each time increment. The second simulation method could question the time duration expected in a particular state, and then jump from one state to the next by knowing the current state duration. This second method is referred to "time-of-flight" or lifetime event simulation since it focuses on a transition *time* rather than a transition *probability* used in the first method. In general, the lifetime event simulation approach is much more computationally efficient, especially for the case of reliable components and systems.

4.4.3.1 Thinning Event Simulation

In the thinning event simulation approach, we focus on the determination of the probability that the simulation will transition from a state to another state within the next time interval. Examples of these transition probabilities (within the next incremental time step) are the probability that a pump fails to operate, the probability that a pressure transducer sends a spurious signal, the probability that a component is repaired, and the probability that an initiating event occurs. The simulation examples that were discussed earlier and the simulation source code in Appendix E all fall into the category thinning event simulation.

In a nuclear power plant, two general types of state transitions are modeled. First, we represent state changes "per demand," where the state transition is modeled via a binomial aleatory model. Second, we represent state changes "per time," where the state transition is modeled via a Poisson aleatory model and the time step is incrementally small. For those cases where the time is small (which it will be during the event simulation) we can write the Poisson equation as (Martz and Waller, 1991)

$$P(t \le T \le t+\Delta t \mid T > t) = \frac{e^{-\lambda t} - e^{-\lambda (t+\Delta t)}}{e^{-\lambda t}} = 1 - e^{-\lambda \Delta t} \qquad (45)$$

where

λ	=	the hazard rate of transitions,
T	=	the time until the occurrence of a state transition,
t	=	the operational time.
Δt	=	a small change in operational time.

Consequently, if we are in state A at time t, given that $T > t$, then the probability that we have a failure in the next Δt is given by Equation 45. Further, if the product $\lambda\Delta t$ is small (less than 0.1), then we can rewrite Equation 45 as

$$P(t \le T \le t + \Delta t \mid T > t) = 1 - e^{-\lambda \Delta t} \approx \lambda \Delta t \quad (46)$$

To perform the simulation, we use a Metropolis technique (Metropolis et al, 1953; Au and Beck, 2001) combined with the probability integral transformation theorem (Bain and Englehardt, 1992). One of the features of probability integral transformation is that we can sample from a cumulative distribution function,[†] F, for a random variable *X* since $x_i = F^{-1}(u_i)$, where u_i is uniformly distributed from 0 to 1 (see pp. 201-202, Bain and Englehardt). In our Equation 46, the cumulative distribution of interest is P(t). Consequently, within the Metropolis routine, P(t) is uniformly distributed and the "candidate" transition (which we will call the state transition criteria) can be found by

$$\lambda = P(t \le T \le t + \Delta t \mid T > t) / \Delta t = Unif / \Delta t \quad (47)$$

where *Unif* is a uniformly distributed random variate from 0 to 1. Thus, if we know the failure rate and the time step, we can call a random number generator at each time step to determine if the transition is allowed or not.

In the case of the binomial model, if the number of trials is one (n = 1) then we have the limiting case which is known as a Bernoulli trial. Since we have only one trial, at most we may have one failure (r = 1). Thus, the binomial model reduces to

$$P(r = 1) = p \quad (48)$$

where *r* is the number of failures and *p* is the probability of a failure per trial. Here, we have a simple state transition criteria for the Metropolis routine, namely p = *Unif*.

Now that we have defined the probability of changing state in the next small time interval for either demand or time-driven processes, a decision-applicable thinning event simulation algorithm will follow the general steps:

[†] The only constraint on the cumulative function is that it must be one-to-one.

1. For a given decision alternative and state, determine a simulated probability value, ρ. The initial state for the decision process is one of the decision advisor inputs that is specified by the user. For example, one decision may be to remain at power. The initial state of the plant in this case would be the full power mode A.

2. Compare the simulated probability (ρ) against the state transition criteria (*p*) using either Equation 46 or Equation 48, as appropriate for the *i'th* time step.

   A. If $\rho > p$, then the new state transition is not permitted. For example, ρ may be 0.61, but the transition criteria for the state may only be 0.05. Consequently, this state *can not* undergo a transition. As a result of this "non event", the algorithm loop counter would increase by one delta time (Δt). The disutility associated with the state would be a function of the state and would be cumulatively tracked as a function of time. For example, if the state is such that the plant is operating and continues operating, the disutility associated with this state is zero. But, if the state represented the plant shutdown awaiting a component repair, then impacts associated with that outcome (perhaps cost) would be accumulated.

      This process would be repeated for each of the desired state transitions as needed. Then, following the state processing (where $\rho > p$), the next simulation iteration would return to step 1 for the $i+1$ time interval.

   B. If $\rho < p$, then the state transition is allowed. Consequently, this state transition *may* cause an upset condition (if currently in a benign state) and may impact one or more performance measures. Therefore, we need to calculate the potential outcome of the transition and store the change in plant state.

If the transition represents a plant upset such as a transient, then the upset immediately impacts attributes such as cost and external attention, as dictated by analysis such as the deterministic economic models discussed earlier. Conversely, positive state transitions may lead to a discontinuation of disutility impacts. For example, if a shutdown plant returns to power, then costs associated with the shutdown would not accumulate further. But, the costs that were recorded would be stored and represent the cost outcome for that decision up to the current time interval.

For those transitions that do not lead directly to a plant upset (for example, a component may fail that does not cause a trip), impacts to the PI attributes should still be estimated. For example, the decision maker may decide to repair an inoperable component while staying at power.

3. The process described by steps 1 and 2 are repeated until all time intervals are evaluated through the end of the overall mission time. For many decision evaluations, the overall mission time is given by the time until the next schedule plant shutdown. Consequently, the total time duration may be measured in months, which may require a large number of time step iterations to complete the analysis.

4. After a specified number of (simulation) iterations, the disutility for each performance measure will be known for each iteration. The expected PI may then be estimated from the stored simulation history.

5. After obtaining the expected PI for the *j'th* decision, the process (steps 1 through 3) are repeated for the *j+1* decision alternative. Upon completion of all the decision alternatives, the uncertainty on each alternative outcome must be determined. The uncertainty determination is accomplished by returning to steps 1 through 4, but, for each new uncertainty iteration, the parameters utilized in the simulation are varied according to their individual probability distribution

functions (i.e., their epistemic uncertainty). For example, the repair probability of a component may be determined from operational data and could be exponentially distributed. For the *k'th* (uncertainty) iteration, the simulation algorithm would pick a value from the exponential repair distribution. In general, all of the parameters utilized in the simulation are uncertain, including state transition rates, component failure probabilities, value tree weights, conditional disutility outcomes, repair times, and cost information.

6. Using the decision rule for expected PI, we determine the preferential decision alternative from the list of alternatives.

While the heuristics we presented are specific to simulating decision processes, they are similar to other simulation methods defined in the literature (Kim and Lee, 1992; Psillakis, 1995; Goel and Ren, 1999; Au and Beck, 2001).

4.4.3.2 Lifetime Event Simulation

In the lifetime event simulation approach, we focus on the determination of times until a transition. Examples of these times are the duration a pump operates, the time until a pressure transducer sends a spurious signal, the time until an inoperable component is repaired, or the time until an initiating event occurs. If we assume that one of these states follow the Poisson conditions described earlier in Section 4.1, then the *time* to a state transition is a random variable and can be represented probabilistically by:

$$P(T < t) = 1 - e^{-\lambda t} \tag{49}$$

where

λ	=	the rate of transitions,
T	=	the time until the occurrence of a state transition,
t	=	the operational time.

If the event of concern is not represented by a homogeneous Poisson process, then the functional form of Equation 49 changes, but the overall simulation technique does not need to be modified.

Using the probability integral transformation theorem (Bain and Englehardt, 1992), the parameter t can be simulated using the expression:

$$t = \ln[\, 1 - P(T < t)] \,/\, (-\lambda) \tag{50}$$

but, 1 - P(T < t) is a uniform random variate from 0 to 1 (we will denote this variate as *Unif*), so

$$t = \ln(Unif\,) \,/\, (-\lambda) \tag{51}$$

Note that *Unif* = 1 - *Unif*, since *Unif* is a uniform random number between 0 and 1. Further, we point out that if $X = -\ln[1 - \mathit{Unif}\,]$, then $X \sim \exp(1)$ (Ripley, 1987).

Now that we have defined the time in a Poisson state (say prior to a plant initiating event such as a transient or tube rupture), a decision-applicable lifetime event simulation algorithm will follow the general steps:

1. For a given decision alternative and state hazard rate, determine a simulated value of t using Equation 51.

2. Compare the simulated time (t) against the mission time (τ).

   A. If $t > \tau$, then the new state arrival is *after* the end of the mission. For example, t might be 3.2 years, but the mission time for the decision may only be 100 days. Consequently, this state *can not* cause an upset condition. As a result of this "non event", the algorithm loop counter would increase by one and the disutility attribute associated with the state would be a function of the state. If the plant was currently at power, the disutility impact is zero (no upset was seen). But, if the plant were shutdown, then the duration of the outage would affect performance measures such as the cost.

      This process would be repeated for each of the desired initiating events. Then, following the processing of the initiators, the next simulation iteration would return to step 1.

   B. If $t < \tau$, then the new state arrival is less than the mission time. Consequently, this state transition *may* cause an upset condition and subsequent impact to one or more disutility attributes. Therefore, we need to calculate the potential outcome of the transition.

If the transition is represented by a plant upset condition such as a transient, then the upset will immediately impact attributes such as cost and external attention, as dictated by analysis such as the deterministic economic models discussed earlier. More complicated upsets (transients leading to long duration outages, loss-of-coolant-accidents, tube ruptures, power interruptions) will have correspondingly higher disutility values for the affected attributes. Note that state transition probabilities following an upset condition may be determined by using either Equation 46 or Equation 48, as appropriate.

For those transitions that do not lead directly to a plant upset (for example, a state of increased plant degradation), impacts to the PI attributes should still be estimated. For example, the decision maker may decide to repair an inoperable component while staying at power. While similar to the plant upset utility calculations, the quantification of these impacts may be of a more subtle nature and, consequently, may be more difficult to estimate. Further, since the impacts are, in some cases, indirect, the uncertainty related to these situations may be larger than those for direct plant upsets.

3. After a specified number of (simulation) iterations, the disutility for each performance measure will be known for each iteration. The expected PI may then be estimated from the stored simulation history.

4. After obtaining the expected PI for the i'th decision, the process (steps 1 and 2) are repeated for the i+1 decision alternative. Upon completion of all the decision alternatives, the uncertainty on each alternative outcome must be determined. The uncertainty determination is accomplished by returning to steps 1 through 3, but for each new uncertainty iteration, the parameters utilized in the simulation are varied according to their individual probability distribution functions (i.e., their epistemic uncertainty). For example, the transition rate of plant transients

may be determined from operational data and could be gamma distributed. For the j'th (uncertainty) iteration, the simulation algorithm would pick a value from the transient gamma distribution. In general, all of the parameters utilized in the simulation are uncertain, including state transition rates, component failure probabilities, value tree weights, conditional disutility outcomes, repair times, and cost information.

5. Using the decision rule for expected PI, we determine the preferential decision alternative from the list of alternatives.

4.5 Explanation of the Decision Calculation

Now that all the individual analysis modules have been described, we will discuss the details of the decision model calculation. Specifically, we will walk through the simulation routine that has been developed for the prototype advisory system. For our decision process simulation, we chose to utilize the lifetime event simulation approach described in the previous section.

Prior to performing the actual analysis, the user, the knowledge base, or a combination of both, will have specified applicable decision alternatives. For an analysis, at least two decision options (e.g., remain at power as-is, repair component at power, shutdown) must be selected. We repeat the calculation described in this section for each decision alternative, making sure to include attributes of the decision that are uniqe.

Step 1 – Initialization

To begin the decision analysis calculation, we first initialize the required variables. A vector array is used to store all applicable variables, thereby simplifying the steps required when performing either sensitivity or uncertainty calculations. Initially, the time variable is set to zero. Also, the outcome for each cardinal-based performance measure is initially set to zero. Each ordinal-based performance measure is set to the first interval (i.e., no impact)

$$cost = 0,\ dose = 0,\ accidents = 1,\ cd = 1,\ attention = 1 \tag{52}$$

Recall that both cost and dose are continuous performance measures, while the other performance measures represent discrete outcomes.

Step 2 – Loop through the user-defined time intervals

The current version of the prototype allows up to three time intervals for each decision. An example of a time interval would be the period of time required to repair a component. Following the repair, the plant operation until the next schedule outage would be considered to be the second time interval.

Step 2a – Simulation during the first time interval, worker impacts

First, we recall the current interval duration from the variable vector array. This duration is defined at T_d.

Next, if the user has specified any worker-related activities such as the work level (in hours) or the average dose level (during the work), then we determine the direct impact on the workers via

$$dose = T_{labor} \times D \tag{53}$$

where T_{labor} is the worker labor time (hr) and D is the average dose rate (Sv/hr) during the activitity.

Labor costs are estimated by

$$cost = T_{labor} \times W_{labor} \tag{54}$$

where T_{labor} is the worker labor time (hr) and W_{labor} is the worker labor average repair cost, which has been noted to be 1,900 euro/hr.

Next, the worker safety must be simulated, using the lifetime event simulation technique. First, we determine the state transition time:

$$t_{acc} = \ln(Unif\,) / (-\lambda_{acc}) \tag{55}$$

where t_{acc} is the transition time; *Unif* is a uniform random number between 0 and 1; and λ_{acc} is the worker incident rate of 1.6E-4/hr (see Table 19). Now, we check to see if the transition time is larger than the time actually worked. If the transition time is larger, then no incident is recorded (i.e., the worker safety variable is left as-is). But, if the transition time is less than the total time to be worked, then a worker incident will occur. Further simulation determines the severity of the incident via the probabilities described in Table 19. For example, the probability of a minor injury, p_{minor}, given an incident, is 2.57E-1. Performing the simulation, we would check to see if:

$$Unif < p_{minor} \tag{56}$$

If the simulation indicates that a minor injury occurs, then the worker safety variable, *accidents*, would be increased to the "minor injury" scale category (which has a value of 2). Otherwise, the simulation continues to see if a major injury, fatality, or multiple fatalities will occur. Any of these types of worker accidents would increment the worker safety variable to the respective scale category.

Step 2b – Simulation during the first time interval, initiating events

Now, we need to determine if an initiating event will cause a transition prior to the end of the current interval. To perform this calculation, the simulation loops through each of the nine initiating event types (see Table 18). The state transition time (for the i'th initiating event) is:

$$t_{IE} = \ln(Unif) / (-\lambda_i) \tag{57}$$

where t_{IE} is the transition time; *Unif* is a uniform random number between 0 and 1; and λ_i is the i'th initiating event frequency (note that the user specifies the plant state for each time interval, which dictates the initiator frequencies). Again, we check to see if the transition time is larger than the current time interval. If the transition time is larger, then no initiating event will occur within this interval. But, if the transition time is less than the current time interval, then the i'th initiating event will occur. Note that in the simulation we assume that the incident will be fixed following an initiating event.

First, let us look at the case where the initiating event occurs. Following an initiating event, we determine the probability of experiencing a core damage, conditional upon the initiating event. Thus, the simulation will test:

$$Unif < P(\text{core damage} \mid i\text{'th initiating event}) \tag{58}$$

If the random number is less than the CCDP, then we have a core damage event (for this one iteration – note that there are typically several thousand iterations performed during the simulation). If we have a core damage event, then we must increment the performance measure variables accordingly. For example, the cost variable will be set to the core damage cost indicated in Table 9, the core damage variable will be set to an index of 2 (1 = no core damage, 2 = core damage), and the external attention will be set to its highest index (a value of 6, indicating a lengthy shutdown).

If there is no core damage, then we still have an initiating event. Thus, we need to account for the impacts associated with the initiating event type. For example, if the initiator is a loss of heat sink, then the plant down time will be down approximately 30 days while the "special" costs will be approximately 2 million euro (see Table 17). Thus, the costs here would be found by:

$$cost_{IE} = T_{down\,time} \times C_{lost\ production} + C_{special} \tag{59}$$

where $T_{downtime}$ is the duration of plant downtime following the initiator, $C_{lost\ production}$ is the cost per unit time due to plant downtime, and $C_{special}$ is the special costs associated with the initiating event.

Step 2c – Simulation during the first time interval, finishing the interval

If we do not have an initiating event, we still must consider decision-specific impacts such as plant down times or reductions in power. For example, if the decision being evaluated is one where the plant is shut down for the entire duration of the time interval, then the costs associated with that decision would be:

$$cost_{decision} = T_{down\,time} \times C_{lost\ production} \tag{60}$$

where $T_{downtime}$ is the duration of plant downtime and $C_{lost\ production}$ is the cost per unit time due to plant downtime.

Step 3 – Simulation during the second and third time intervals

The calculations described in Step 2 are repeated for the next two time intervals. Note though that during the next time interval, the plant state may change. For example, the plant may have been shut down in the first interval and may be at power during the second interval. Thus, many variables such as the initiating event frequencies and the CCDPs are automatically adjusted to account for the particular plant state.

Further, one should realize that the performance measure variables are accumulated through each of the time periods. By this we mean that the overall total cost after the three time intervals is:

$$cost_{total} = cost_{duration\,1} + cost_{duration\,2} + cost_{duration\,3} \tag{61}$$

The dose, attention, core damage, and worker safety measures will also be accumulated over the disjoint time periods.

Step 4 – Determine the PI outcomes

Following the three time period calculations, we need to determine the PI outcome for each of the five performance measures. This calculation, for cost, is performed by:

$$PI_{cost} = w_{cost} \times u(cost_{total})_{cost} \tag{62}$$

where w_{cost} is the cost performance measure weight and $u(\ )_{cost}$ is the disutility at a value of $cost_{total}$ (see Figure 31).

For the dose calculation, we would have:

$$PI_{dose} = w_{dose} \times u(dose_{total})_{dose} \tag{63}$$

while the other performance measure would have similar expressions.

The total PI for the decision is then the summation of the individual PI values, or:

$$PI_{decision} = PI_{cost} + PI_{dose} + PI_{accidents} + PI_{core\ damage} + PI_{attention} \tag{64}$$

Step 5 – Determine the PI statistics

Since the simulation will be repeated several thousand times (or iterations), we will need to calculate PI statistics such as the expected PI:

$$E[PI_{decision}] = \frac{\sum_{i=1}^{n} PI_i}{n} \tag{65}$$

where PI_i is the PI from the *i'th* iteration and *n* is the number of iterations. If we were performing an uncertainty analysis, the individual variables stored in the vector array (described in Step 1) would be randomly modified via Monte Carlo sampling. Consequently, the simulation process would have to be repeated for the desired number of uncertainty analysis iterations. Thus, the overall calculation time for simulation may be significant, which is one of the typical drawbacks to this type of calculation.

4.6 Summary

We have now outlined the decision *analysis* for the incident management prototype. Within this analysis framework, we described both the deterministic and aleatory models required to assist in the calculation of preferential decision options. In general, the calculation of preference between decision alternatives focuses on the expected PI, or:

$$E[PI] = \int_0^{\infty} \left(\sum_{i=1}^{5} w_i \, u(x_i \,|\, y_i) \right) \pi(\bar{\mathbf{x}}_i \,|\, \bar{\mathbf{y}}_i) \; d\bar{\mathbf{x}}_i \tag{66}$$

Since we are dealing with disutilities [$u(x_i \mid y_i)$] for the prototype system, we seek to have decision options with low E[PI]. But, as part of the expected value calculation, we need to incorporate epistemic uncertainties [$\pi(x|y)$] into the analysis. Thus, we denoted the types of sources and magnitudes of uncertainties for the items such as economics, judgement, plant operation, and the safety assessment.

During the analysis discussion, we noted that we rely on two types of models, deterministic and aleatory. Within the deterministic framework, we focused on the models of plant operation and economics. For plant operations, the initiating events were used in two ways, first to determine the probability of getting to a decision-specific state (say a transient) and then second to determine the outcome of that state. Our decision maker has specified the likely impacts that would be realized following any one of the nine categories of initiating events (Table 17). For the economics modeling, we utilized the work done by Burke, Aldrich, and Rasmussen (1984) in their report on nuclear power plant costs. For example, from this work we were able to determine hourly repair costs, worker injury costs, and core damage costs.

The aleatory modeling focused mainly on plant operations and the safety assessment. Within the context of plant operations, we have modeled plant upset conditions as leading to impacts in the PI performance measures such as core damage, cost, and external attention. Further, we determined models for industrial accidents and radiological dose, both of which affect workers. For example, the industrial accident model utilized worker accident rates collected for the U.S. Department of Energy CAIRS database. This database includes a variety of accident information spanning the years of 1995 to 2000.

We then went on to describe the need and approach to simulation-based calculations for decision processes. The outcome represented in the decision model must exhibit an appropriate level of detail regarding potential outcomes following a decision. For example, if a plant experiences more than one transient due to a decision to remain at power, the PI for the outcome *must* reflect the multiple outages. If each outage following a plant trip results in lost power generation, the cost impact could be quite large. Further, if a plant experiences multiple trips in a short period of time, other impacts like external attention may be increased due to the sensitive nature of initiating events at nuclear power plants. We also compared static model calculations against simulation models, and illustrated the areas where simulation proves superior.

The goal of the decision-process simulation is to determine potential outcomes, determine their individual impacts, weight the impacts by their respective likelihood, and then calculate an overall E[PI] for each decision alternative. Then, the decision alternative with the lowest E[PI] is considered to be the preferred option. In order to implement the event simulation, we provide the mathematics needed for both a “thinning event” simulation and a “lifetime event” simulation. Thinning event simulation questions the state transition probability at each incremental time step. Lifetime event simulation questions the time duration expected in a particular state, and then proceeds from one state to the next by knowing the current state duration. In general, the lifetime event simulation approach is much more computationally efficient, especially for the case of reliable components and systems. Consequently, this is the method we implemented for the simulation module. Further, we described, step-by-step, the details of the calculation involved in simulating a decision.

We have now combined the modeling techniques and analysis methods into a prototype incident management system. This system is described, and utilized, in the next section. We used the prototype system to evaluate actual incidents that occurred at two of the decision maker's plants. For this discussion, we will explain the context of the incident, the modeling required when using the prototype, and the results that are derived from Version 1 of the prototype[†].

[†] The culmination of approximately 4,000 lines of PHP code has gone into Version 1 of the prototype.

"Facts do not cease to exist because they are ignored." — Aldous Huxley

5 Application Examples for Incident Decision Making

Recall that it is the goal of the incident advisor prototype to facilitate selection of a preferential decision alternative from available options and provide technical justification for the basis of the decision. In previous sections, we defined the salient features of the decision model and presented methods of solving such models. In this section, we demonstrate the methodology, and the advisor prototype, through two case studies. The first case study deals with decisions concerning a degraded, but not failed, component, namely a small leak in a steam generator tube. The second case study investigates the decision process for a failed component, a pressure transducer attached to the primary coolant system. We will provide an overview of each incident, denote the inputs to the decision model, and then analyze the decision model to determine a preferential option. In both cases we will demonstrate use of the prototypical decision advisor software developed as part of this work. But first, let us describe the prototype system.

5.1 Structure of the Decision Advisor Prototype

In order to bring together the decision modeling and analysis heuristics, we have developed a decision advisor prototype. The goal of the prototype is to assist nuclear power plant personnel in their response to incidents through implementation of the methods described in this report. During the development, we focused on four parts:

1. A primary controlling module to collect incident information and subsequently determine the decision model.
2. Preference models representing the decision maker's beliefs (via a value tree and associated disutilities).
3. Supporting analysis modules (for example, the PRA, economic models, worker safety models).
4. An analysis module to solve the decision process.

While these four modules make up the structure of the prototype, we have developed a list of salient features. These features of the prototype include:

Elicit the General Context of Incident. The prototype should be able to address incidents involving either degraded SSCs in addition to initiator-type situations. Further, relevant information such as the plant boundary conditions and other event specific attributes driving the incident must be known. Once the relevant initial information is supplied to the prototype, it should then be able to guide the analysis, with the assistance of the user, to the completed decision results.

Construction and Solution of the Primary Decision Model. The prototype should help the analyst construct the decision model that will be used to determine an incident strategy. This function should include both the ability to construct new models "on-the-fly" and the storage/search capacity to recall prior evaluation that may be similar. It is envisioned that the primary decision model that will drive the analysis is an influence diagram. The attributes defined in the value tree will be present in the decision model, primarily via the "outcome" node that represents the overall disutility. Consequently, the prototype must utilize the PRA model, an economics module, worker risk information, and estimation of external attention in order to provide a solution for the incident in question.

Provide Modeling Information. The prototype should be able to provide information (either through a knowledge base or an algorithm) relevant to the decision making process. This information will include items such as overall model structure, decision alternatives (e.g., continue operating, shutdown, repair at power), probabilistic influences (e.g., chance nodes, uncertainties), and outcome information (e.g., economic impacts, non-economic impacts, or disutility). Since some of the analysis may be internal to the prototype, it is important that the user be informed as much as possible as to the "inner workings" of the analysis.

Assist in the Quality Assurance of the Decision Model. The prototype should assist in checking the model and results of the decision process. It is believed that many automated checks could be implemented to prevent modeling errors. For example, the prototype could check to ensure that entered or calculated probabilities fall between 0 and 1. Also, checks could be provided to canvas the outcomes such as economic costs to make sure that they fall within feasible ranges.

Provide Results and Sensitivity Metrics. The prototype should yield applicable results (either textual or graphical) with minimal burden on the user. This reporting capability should be tailored to the individual requiring the analysis results. Results should include information such as "tornado diagrams" (i.e., graphical sensitivity reports) and the multi-way sensitivity plots. Also, information specific to the uncertainties on the calculated decision results should be provided in order to give the user an understanding of the quantitative aspects found in the results. The prototype should be able to provide information (beyond the normal results) that lends insights into the decision models and the drivers behind the models. Since it is desirable to abstract and automate as much of the decision modeling as possible, the prototype could yield results that are not immediately understood by the analysts. Consequently, there is a real need to be able to have the prototype communicate the "how" and "why" of any analysis.

We have taken these desired features and the knowledge of the four main modules of the prototype to construct an information and analysis flow framework. This general framework of the prototype is shown in Figure 55, where a total of five stages are represented. Each stage embodies a unique portion of the overall process of determining a preferential decision option from a list of alternatives. We will discuss each tier in turn.

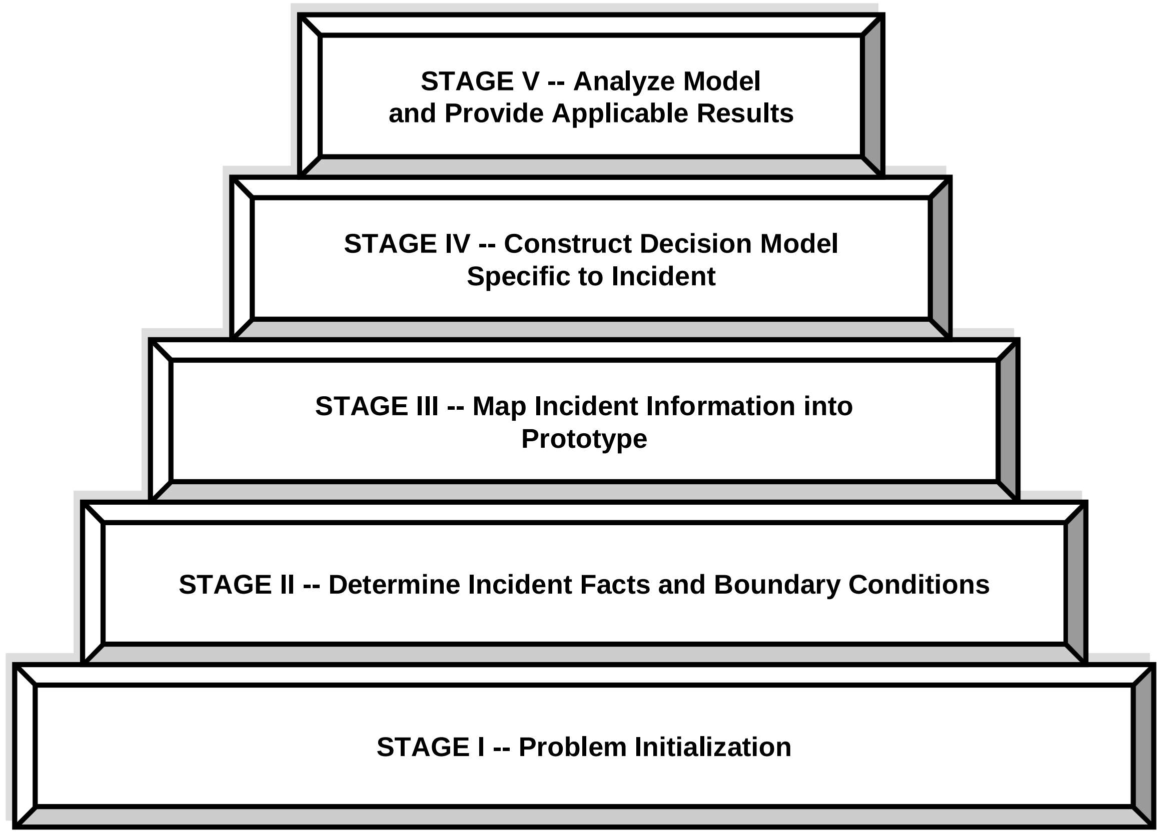

Figure 55. Framework tiers embodied in the decision advisor prototype.

Stage I of the prototype represents the initialization of the decision process. The current version (version 1) of the prototype uses a web-centric approach where the prototype runs off an web server.[†] During the initialization process, the user must first log into the advisor system. Then, the user is presented with initial information, namely the value tree, performance measure weights, and associated disutilities. This information is available for review but can not be modified (in this version of the prototype) by the user. An example of the graphical user interface for the prototype is shown in Figure 56. Note that the prototype is multi-lingual, where the language can be changed by selecting from a list of available options.

[†] The programming language used to develop the prototype is PHP, version 4 (www.php.net).

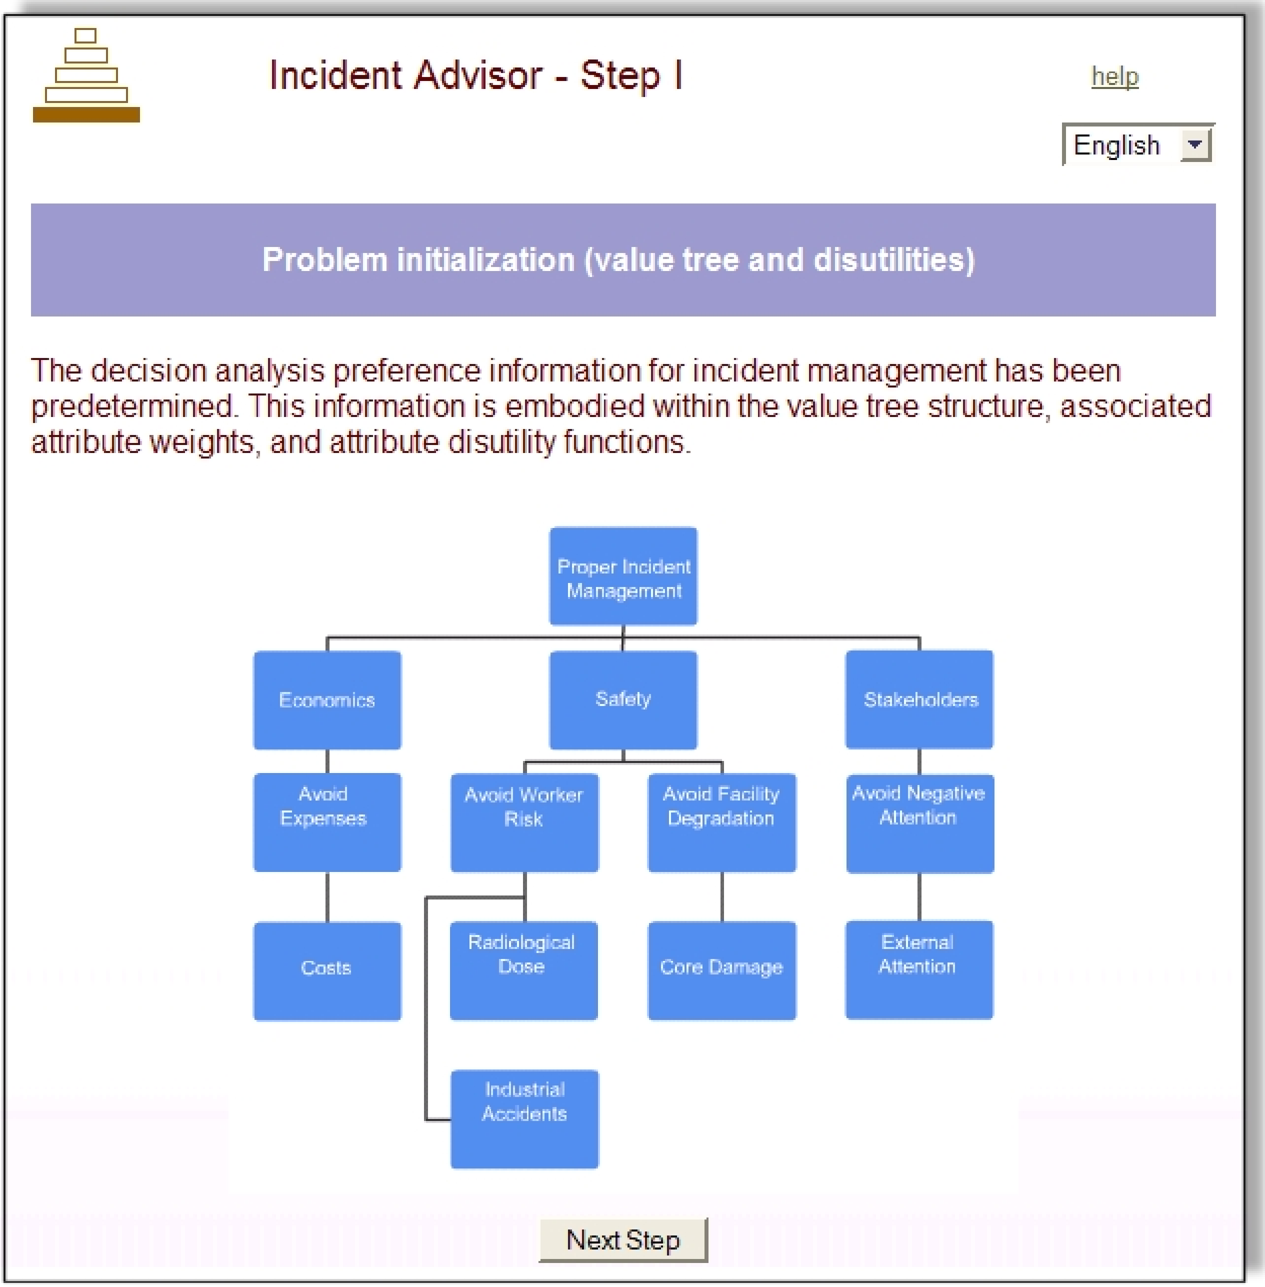

Figure 56. Example of the user interface for the decision advisor prototype.

We have discussed the knowledge base that is embodied in the prototype. Using this knowledge base, the prototype queries the user for relevant information applicable to the incident. In later stages, the prototype will adjust the decision model based upon the user input; allow modifications to both models and associated data as the analysis is constructed; solve the model; and then provide results. But, in Stage II, the user must provide incident-specific facts and related boundary conditions. For example, types of data that must be supplied by the user at this stage include:

- The incident name and description
- The type of incident (component or initiator related)
- The current reactor state
- The time until the next scheduled outage
- Impacts to plant operations through component degradations

Often times the user may not know directly how a component degradation affects plant operations. For example, in the case of our failed pressure transducer, the degradation affects the likelihood of tripping the plant from full power. Unfortunately, this component is *not* modeled in the decision maker's PRA. When developing a PRA, the analysts determine which, of the tens of thousands of components, to model and which to leave out. If a component is modeled, it is represented by one or more "basic events" in the PRA. For example, a diesel generator may be modeled by fails-to-run, fails-to-start, and maintenance outage events. If a diesel generator is degraded, the user may select this component directly from the list that is provided in Stage II. But, if a component is not represented in the PRA, one may still determine its impact on the PRA by use of "surrogate components."

A surrogate component is needed only when a component is not in the PRA logic models. Current nuclear power plants have, as a rough estimate, over 50,000 components. But, a modern PRA typically models between 1,000-4,000 basic events. Consequently, the number of components directly in the PRA is substantially less than that in the plant itself. Consequently, we also ask the user if degraded components potentially impact either (a) initiating events or (b) plant safety systems. If the user answers affirmative, then the user must subsequently specify these impacts.

Within Stage II, the user will provide a variety of facts and boundary conditions specific to the incident. Some of this information will cause the prototype to ask the user for additional information (for example, if a component degradation impacts initiating events, the user will be queried for this information) in later stages.

In Stage III, the details of the incident are entered into the analysis data stream. A variety of data that resides in the prototype knowledge base will be revealed as nominal or “base case” values – the user may override this information within this stage. In the case of a failed pressure transducer, the component degradation is not in the normal PRA. But, this component impacts transient initiating events. A transducer failure makes it more likely to have an upset since any one of the three remaining transducers can cause a trip, in other words the first failure is a “fails safe” type that changes the system from a 2-of-4 to a 1-of-3 system.

Once the user indicates that a failure can impact entities such as initiators, they will be asked to note the value of the impact. The data entry screen looks like that shown in Figure 57. The values that appear in this figure are default frequencies for the initiating events listed previously in Table 18 for each of the major plant states. To modify one of these values (for example the primary transient category), the user clicks on the applicable data entry. This process of indicating impacts to the decision model is continued through the nominal configuration and any potential decision alternatives.

The knowledge base contains a selection of potential decision alternatives, options like continuing as-is, shutting the plant down to fix the problem, repairing the problem at power, etc. Incident conditions that are tied to a decision (through the knowledge base) will cause that decision to automatically be used in the analysis. An example of this case is for a failure of the diesel generator, where it is tied to the option “repair at power” since this component is a standby safety system that may be modified while the plant operates. Conversely, an incident involving the accumulator† could not be repaired at power since it is an integral part of the primary coolant system – thus it would not have an association with that decision node in the knowledge base. After the user denotes the final list of potential decision alternatives, additional details for each option must be provided.

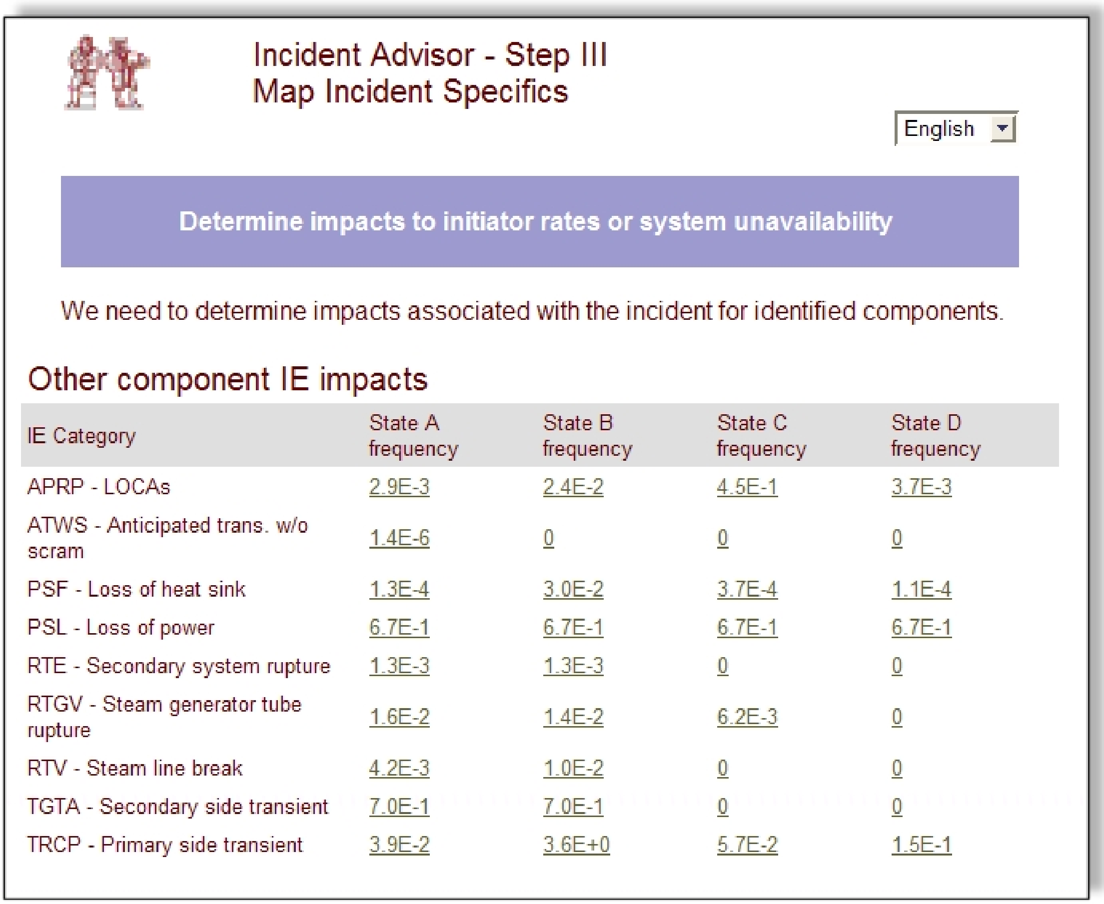

Figure 57. Example of the data entry interface of the decision advisor prototype for impacts to initiating events.

---

† The accumulator is a small vessel, partially filled with water, which helps to control the pressure inside the reactor coolant loops.

Version 1 of the prototype allows a total of three plant state changes following the implementation of a decision. Most decisions will require fewer than three states to represent the decision outcome. Examples of the states associated with decision alternatives are (assuming the plant is originally in the full-power state A1):

- If the decision is to remain as-is, then plant will (nominally) remain in A1 until the next planned outage. The decision model for this alternative will then assume the plant is in state A1 until it leaves this state due to a stochastic event such as a transient.

- If the decision is to repair the component while the plant remains at power, the decision model will use two states. First, the plant will be in state A1, but only for the time it takes to repair the component. For this period, the prototype will ask the user for the duration, plant degradation values (to either initiator or system failure likelihood), and the worker hours needed to fix the inoperable component (in order to determine the industrial safety impacts). For the second time period, the plant will stay in state A1, but the plant is repaired (if a successful repair) so the plant status will return to the nominal, fully functional, state.

- If the decision is to shut the plant down to fix the problem, the decision model will use two states. First, the plant will go to a shutdown state (for example, state A3, where the reactor is subcritical, but the plant is in a “hot shutdown”) for a short duration. The user must specify the actual state, plant impact while in this state, and the duration of time in this state. Following the repair, the plant will return to the full-power state A1.

Once the relevant decision information is entered into the prototype, the user enters Stage IV. Here, the prototype will construct the decision model that will be analyzed to determine the preferential decision option. Prior to the actual analysis, the prototype displays the information that has been collected by the user in order to offer a final check of the decision model inputs.

Stage IV of the prototype is where the decision model is constructed (internally) by the prototype. The current process that takes place within this stage is to quickly evaluate the generic influence diagram that was discussed in Section 3.2 via a static, sequence-based "roll-back" calculation (Clemen, 1996). This numerical calculation is effectively the same as that provided by the standard decision tree approach. While this calculation may not be as robust as that from simulation, it is nonetheless much faster. The prototype spends approximately three to five seconds solving the static model, including a low and high sensitivity calculation for the major model variables.

Once the static results are shown on the screen, the user has officially entered Stage V of the decision advisor prototype. Stage V is the final step in the analysis process, whereby the user is allowed to explore the analysis results and, if necessary, instruct the prototype to evaluate the model using simulation. At a minimum though, the user is shown the decision alternative ranked by preference along with the numerical score (calculated via Equation 5). Also, the user is allowed to view sensitivity calculations and perform a Monte Carlo uncertainty analysis. Examples of these screens will be shown during the case study discussion.

5.2 Case Study I – A Leaking Steam Generator Tube

5.2.1 An Overview of the Problem

In a pressurized water reactor (PWR) nuclear power plant, coolant leaves the reactor and enters one of the steam generators. In the steam generators, heat is transferred from the primary coolant to the secondary. A typical steam generator is of a "U-tube" type with the primary fluid inside the U-tube and the secondary fluid flowing over the outside of the tubes. Of concern with this system is the potential for leaks or ruptures of the one or more U-tubes. Leaks of primary coolant to the secondary will cause an increase in the secondary radioactivity. Larger leaks or ruptures may cause a challenge to the ability to cool the reactor and provides a direct path from the reactor to systems outside the reactor containment structure.

The incident modeled for Case Study I was that of a leaking steam generator tube. At a typical PWR, small leaks of the steam generator can be mitigated by operation of the chemical and volume control system (CVCS). Note that the definition of small leaks (in this report) include incidents where the primary-to-secondary rate is less than 20 l/hr. Thus, if a small steam generator leak occurs and the CVCS functions, the inventories in the primary and secondary remain stable, but the secondary radioactivity level will gradually increase. Further, for small leaks, the reactor will not automatically trip, but if desired, the operators may reduce power and manually trip the reactor. If the leak rate were above 20 l/hr the plant would have to proceed to a shutdown situation as described by its technical specifications.

5.2.2 Key Inputs into the Decision Model

If the leak is small, the incident may last over weeks or even months. Very small leaks from the primary to secondary are a normal part of operating any nuclear power plant. So, the decision-maker must evaluate, almost on a constant basis, the basis for future operation given the current conditions. One may want to look at this case study as evaluating the same decisions each day until the problem (i.e., the leak) is repaired. Or, we may need to revisit the decision making any time additional information relevant to the decision is known.

The PRA associated with the prototype has an initiating event of steam generator tube ruptures. But, these tube ruptures are much larger leaks than that represented by our incident. Unless the leaking tube suddenly ruptures, the incident is more like an inoperable component that may cause a plant transient rather than a full-fledged tube rupture. Consequently, to model a leak in the steam generator tube rupture, we used the prototype to adjust the secondary side transient frequency slightly upward by 10%. This adjustment value is an assumption based upon engineering judgement since we did not have a reliability model that would translate tube leaks into transient upsets.

Since small leaks from the steam generator tubes are not currently integrated into the prototype knowledge base, we modeled Case Study I by indicating that the incident "impacts initiating events." We also assumed that when the leak was found, the plant was in a full-power mode, State A1, and the next schedule outage was 100 days into the future.

5.2.3 Analysis of the Decision Model

For each incident, the prototype offers a list of possible decision alternatives. But, since small leaks from a steam generator are not explicitly modeled in the knowledge base, the prototype will not offer incident-specific advice related to the potential alternatives. Instead, we are allowed to select appropriate alternatives from the master list. Doing this, we picked the following decisions:

1. Continue running with the primary system as-is (with one pressure transducer failed safe and three working). Here we maintain plant as is and use normal makeup (i.e., CVCS) to mitigate the leak. The secondary will have an increase in radioactivity due to the leak, which we can model by specifying an increased level of radiation work level.

2. Reduce Power. Here we are going to bring the plant to a lower power level (e.g., 80-90%) in order to mitigate the leak rate. But, the secondary coolant system will continue to have an increase in radioactivity due to the leak.

3. Shut the plant down and perform the repair. Here we must bring to plant to a cold shutdown state in order to repair the leaking tube. Repair could consist of resleeving or plugging the suspect tube(s). It was assumed that the leaking tube could be plugged within 96 hours.

5.2.4 Case Study I Results

The prototype analysis results screen is shown in Figure 58. As can be seen, the preferential decision is to continue operation of the plant as-is. The calculated PI for this alternative is about 60% lower than the next closest option. Two of the major reasons why the "continue as-is" option is, in this case, preferential are: (1) secondary side transients, if they occur, are relatively benign and have a low CCDP and (2) the increase in the secondary side transient frequency was estimated to be fairly small (10%).

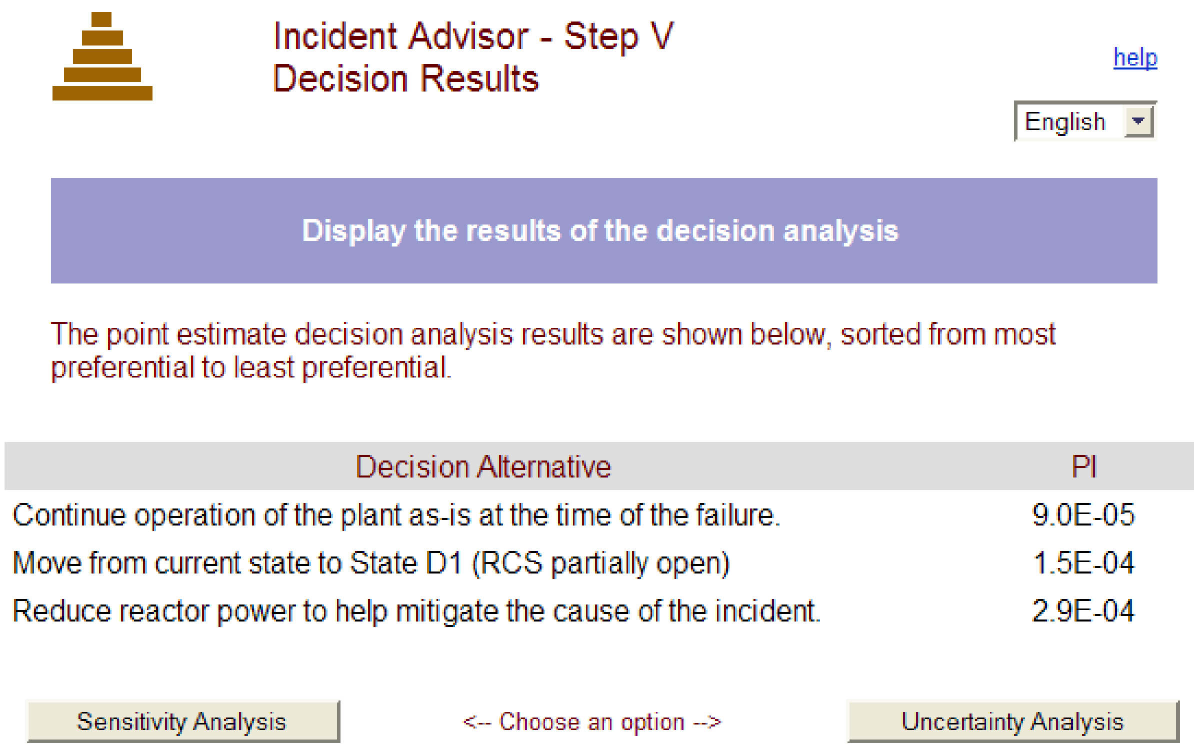

Figure 58. Point estimate results screen from the decision advisor prototype for Case Study I.

The prototype has been programmed to perform sensitivity analyses on all of the relevant variables that enter into the decision model. For each variable, we recalculate the overall expected PI for both a low and high value. Then, we determine which variables have the largest impact (from low to high expected PI). Of these, the prototype plots the top 15 variables as a “tornado plot.” This plot is shown for Case Study I in Figure 59.

Evaluating the tornado diagram, we can see that a couple secondary side parameters are important to the analysis. But, one should also notice that other upset conditions, for example loss of power, might be more important. This importance implies that events like loss of power, while occurring less frequently than secondary transients, have a larger impact (in decision space) than more benign events.

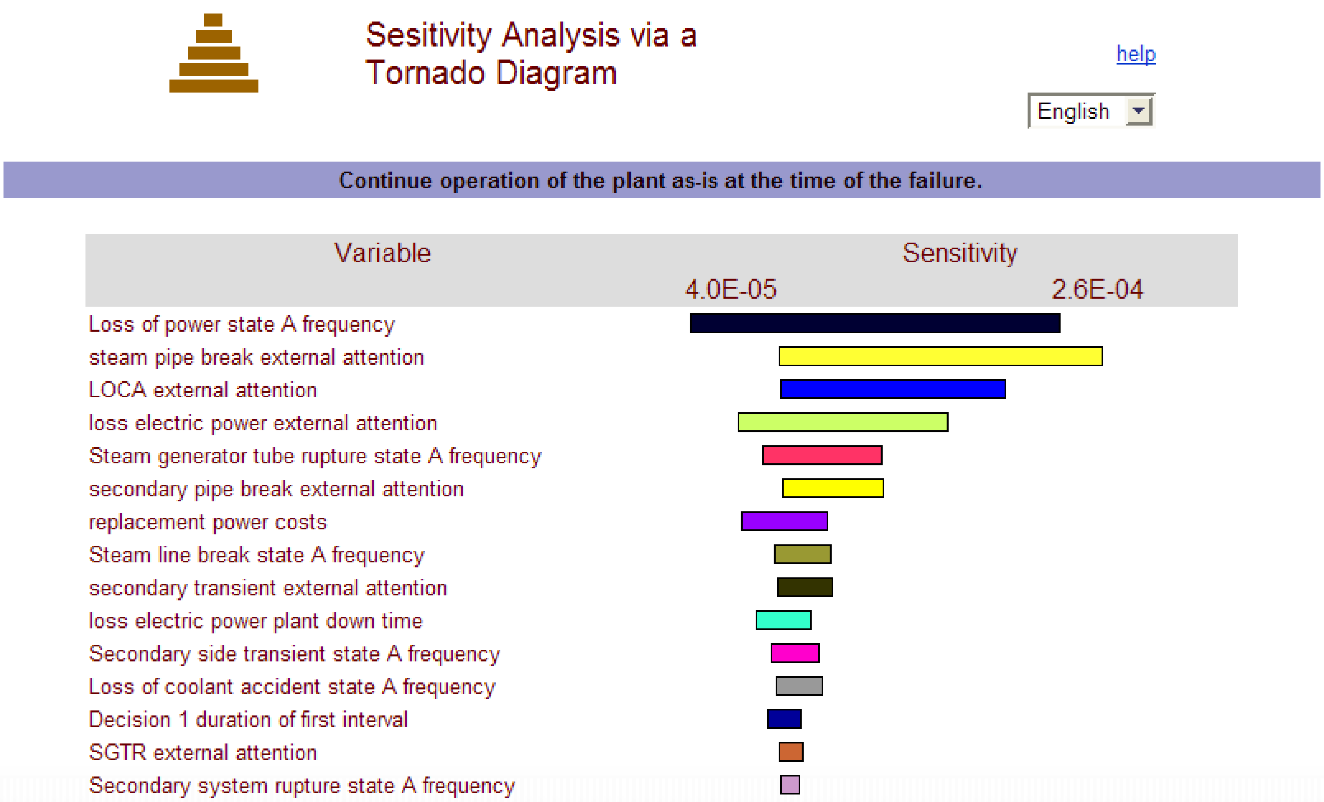

Figure 59. Sensitivity results (via a tornado diagram) for Case Study I.

The prototype has also been programmed to perform an uncertainty analyses on all of the relevant variables that enter into the decision model. For each variable, we recalculate the overall expected PI, for all decisions, using their epistemic uncertainties. Then, for each uncertainty iteration, the prototype determines which decision is most preferred, which one is second, which one is third, etc. The positional preference order is then tabulated over all the iterations to determine the fraction of time the *i'th* decision is preferred over the other decisions. These results are then plotted. An example of the uncertainty results is shown for Case Study I in Figure 60. In this figure, we see that the preferred decision (continue operation) from the point estimate results is preferred in approximately 80% of the iteration calculations. The other two options are similar in that they trade second and third place by about the same fraction.

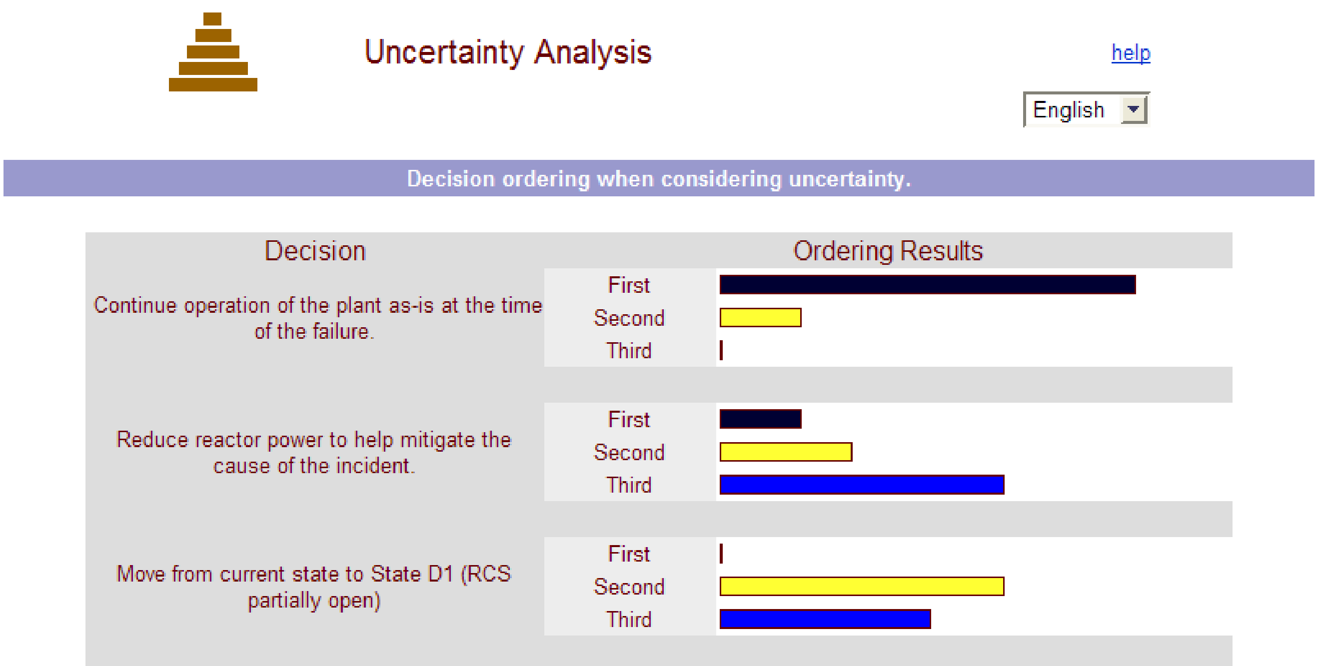

Figure 60. Uncertainty results for Case Study I.

5.3 Case Study II – A Failed Primary Coolant Loop Pressure Transducer

5.3.1 An Overview of the Problem

In a nuclear power plant, portions of the primary system are instrumented to provide process signals such as pressure and temperature. The instrumentation that provides these signals is typically in harsh environments (e.g., high temperature and humidity) and may be difficult to repair. Since it is desirable to have reliable signal outputs, this instrumentation is arranged so that multiple trains can provide a reading. For example, a temperature signal may be constructed such that three instruments provide a temperature measurement at approximately the same location.

While a plant has redundancy built into the instrumentation equipment, failure of one of these signal devices causes concern to the plant. Not only is the reliability of that signal degraded, but the plant is now susceptible to spurious signals since most instrumentation is designed to fail-safe. For example, if a pressure signal has four redundant trains, where any two can initiate safety injection, failure of one train leaves the system in a "one-out-of-three" arrangement.

The second case study represents the incident where the plant experiences an inoperable pressure transducer in the primary coolant system. If one additional transducer emits a spurious signal, the plant will have a transient.

5.3.2 Key Inputs into the Decision Model

The knowledge base behind the prototype has been augmented with information specific to the pressure transducers. First, a reliability model, using fault tree minimal cut sets, has been integrated into the "systems analysis" module. This module will now allow us to determine the likelihood of the plant tripping due to a spurious signal from the pressure transducer system. In other words, we are extending (slightly) the PRA, but this modification now becomes a permanent addition to the prototype. Consequently, in the future, if the decision maker has an incident involving a pressure transducer alone, or in conjunction with other failed components, the knowledge base will provide a solution to initiator frequency impacts. The second addition to the knowledge base was the inclusion of influences on nodes within the database. Since a failed pressure transducers may be repaired at power, a relationship was identified between all four transducers and the "repair at power" decision node. Alternatively, one could shut the plant down and fix the problem, or since there exists three backup transducers, one could continue to operate the plant as-is. Thus, these additional influencing factors were incorporated into the knowledge base.

Since the prototype knowledge base handles much of the analysis, all we need to do is identify the failed pressure transducer from the list of modeled components. Following this identification, the prototype will identify three candidate decision alternatives (repair at power, remain as-is, shutdown and repair), but we are allowed to add additional alternatives as necessary. For each alternative, the prototype indicates the states and transition times that may be expected. For example, remaining at power would (nominally) only require one plant state (the current one) and the duration would be until the time of the next schedule shut down (which was specified earlier). Of course, the user can override the default values.

The prototype also allows the inclusion of repair-caused plant upsets and the possibility that a repair is performed incorrectly. The default is that the decision model will utilize these parameters, but the user may elect not to select these. For this case study, we did allow the prototype to factor in these potential impacts for the decision model.

5.3.3 Analysis of the Decision Model

As discussed, the prototype offers the following possible alternatives when a pressure transducer fails:

1. Continue running with the primary system as-is (with one pressure transducer failed safe and three working).

2. Perform a temporary repair while at power (i.e., on-line repair)

3. Shut the plant down and perform the repair.

The first alternative is optimal if we only considered cost (and there are no other faults in the other pressure transducers). But, operating the plant with a degraded safety system opens the potential for an inadvertent plant upset. In this case, we would experience both an economic loss due to the outage and could have negative attention from the public and/or regulator.

The second alternative, perform an online repair, will restore the system to an as-good-as-new status. As part of this strategy, we may consider the likelihood that the repair will be performed improperly. If the repair were successful, we would only sustain the costs of the repair and could run the plant without stopping. But, there is a small probability that the repair personnel will cause an inadvertent transient (while doing the repair on-line), thereby shutting the plant down. This human error would cause an economic loss, as well as increase the potential for negative attention.

The third alternative would be to shut the plant down and repair the component. The decision maker would avoid negative attention, but would face a guaranteed economic impact due to the lost energy production. Fortunately though, the decision advisor prototype allows us to consider tradeoffs between performance measures such as cost, safety, and adverse attention. It was assumed that repair would take 48 hours.

For this case study, we assumed that there were 100 days until the next schedule shutdown. The plant state at the time of the incident was full power, State A1. We did not add any decision alternatives to the list of three provided initially by the prototype. Also, we allowed for the possibility of human error via a repair-caused trip and an unsuccessful repair.

5.3.4 Case Study II Results

The prototype analysis results screen is shown in Figure 61. As can be seen, the preferential decision is to continue operation of the plant as-is. The point estimate PI for this alternative is slightly lower than the other two options. But, three reasons why the “continue as-is” option is, in this case, preferential are:

1. Primary side transients, if they occur, are relatively benign and have a fairly low CCDP.

2. The downtime impact is small (see Table 13).

3. Pressure transducers are relatively reliable components, giving operational margin when operating with only three available.

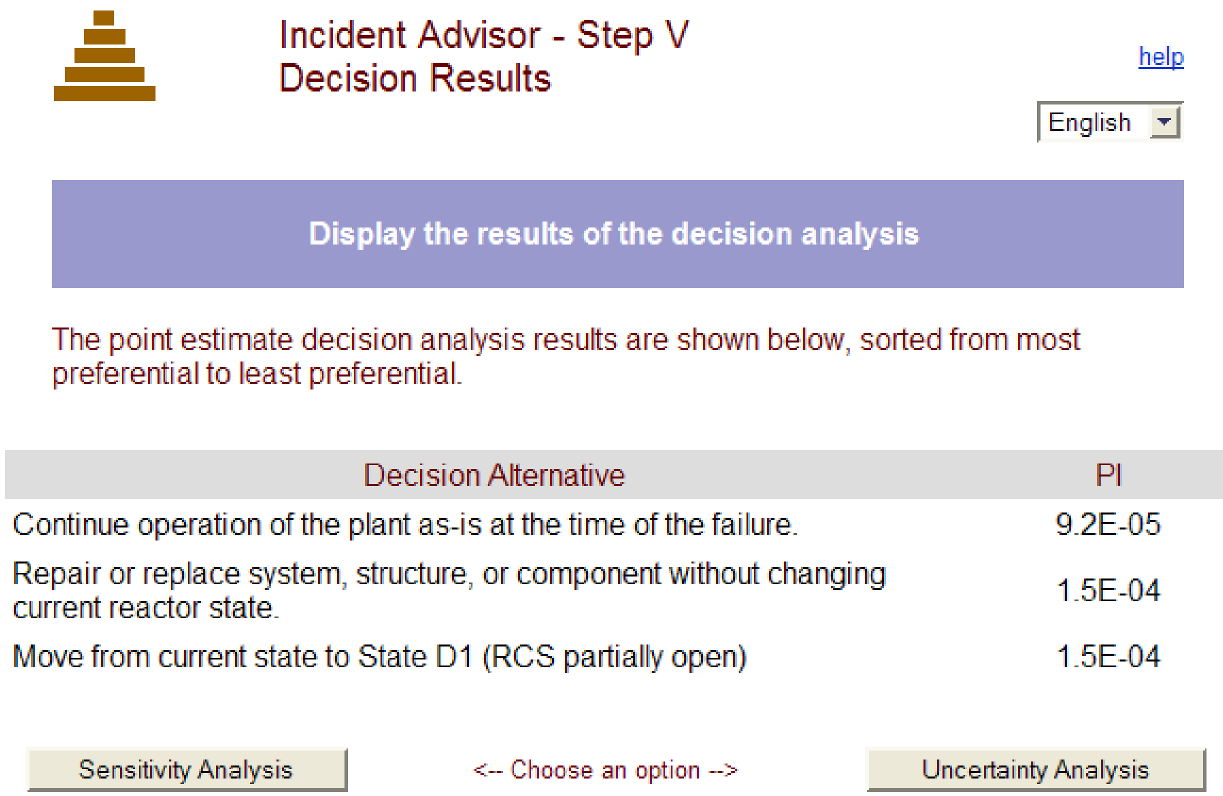

Figure 61. Initial results screen from the decision advisor prototype for Case Study II.

Again, a sensitivity plot was constructed for this case study, whereby we determine the top 15 variables. This plot is shown for Case Study II in Figure 62.

Evaluating the tornado diagram, we can see that the primary side transient frequency is not as important as other variables. Since the pressure transducer system only partially impacts the primary transient value, other parameters associated with larger impacts tend to be more important. For example, note that secondary side transients are more likely than primary transients (see Table 18), even during the degraded state. But, these two transients have similar outcomes with regard to impacts on PI. Thus, parameters such as secondary transients would tend to dominate over primary transients since the tornado diagram represents the decision to continue operation as-is.

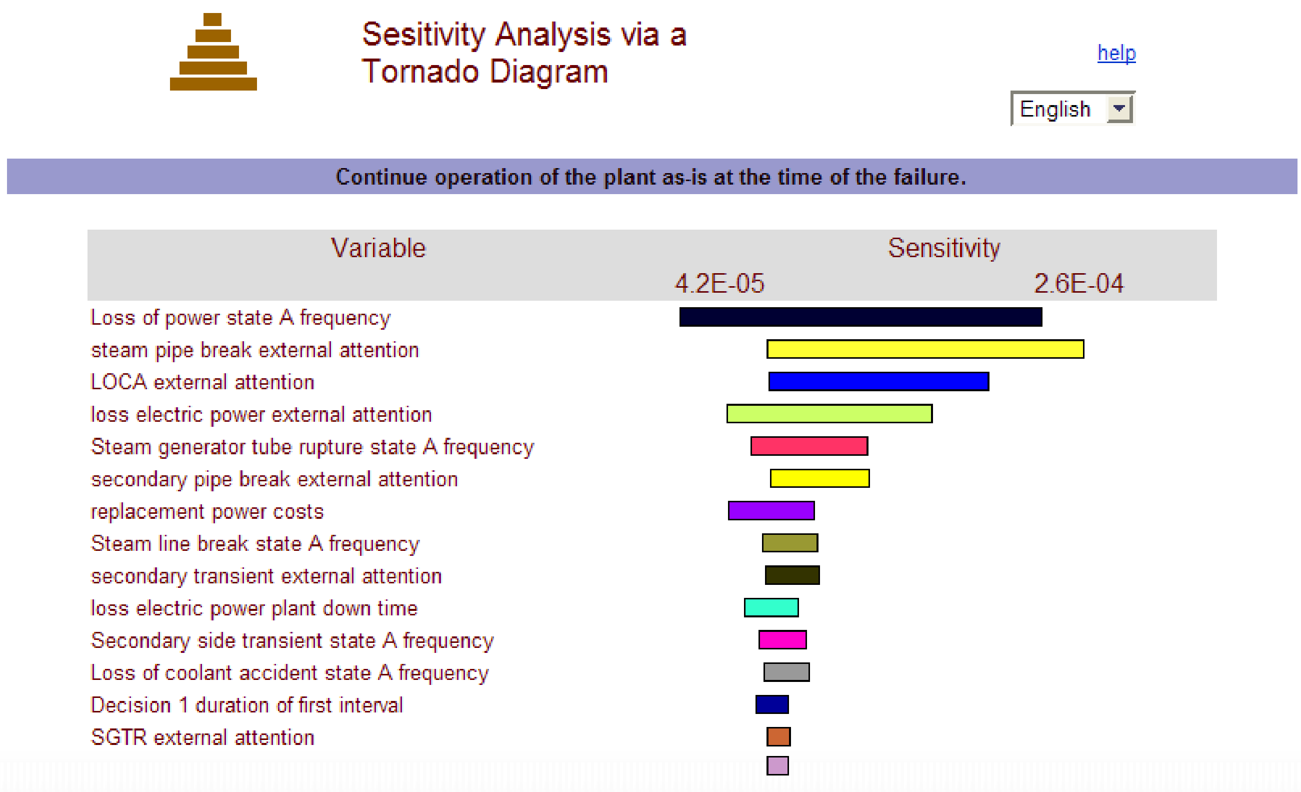

Figure 62. Sensitivity results (via a tornado diagram) for Case Study II.

An example of the uncertainty results is shown for Case Study I in Figure 63. In this figure, we see that the decision to continue operation is preferred in almost 100% of the iteration calculations. While this “dominance” result may be surprising, it is the same as the case described in Section 4.1.3. Repairing the transducer at power generally falls second (in the list of preferential decisions) while shutting the plant down is typically third.

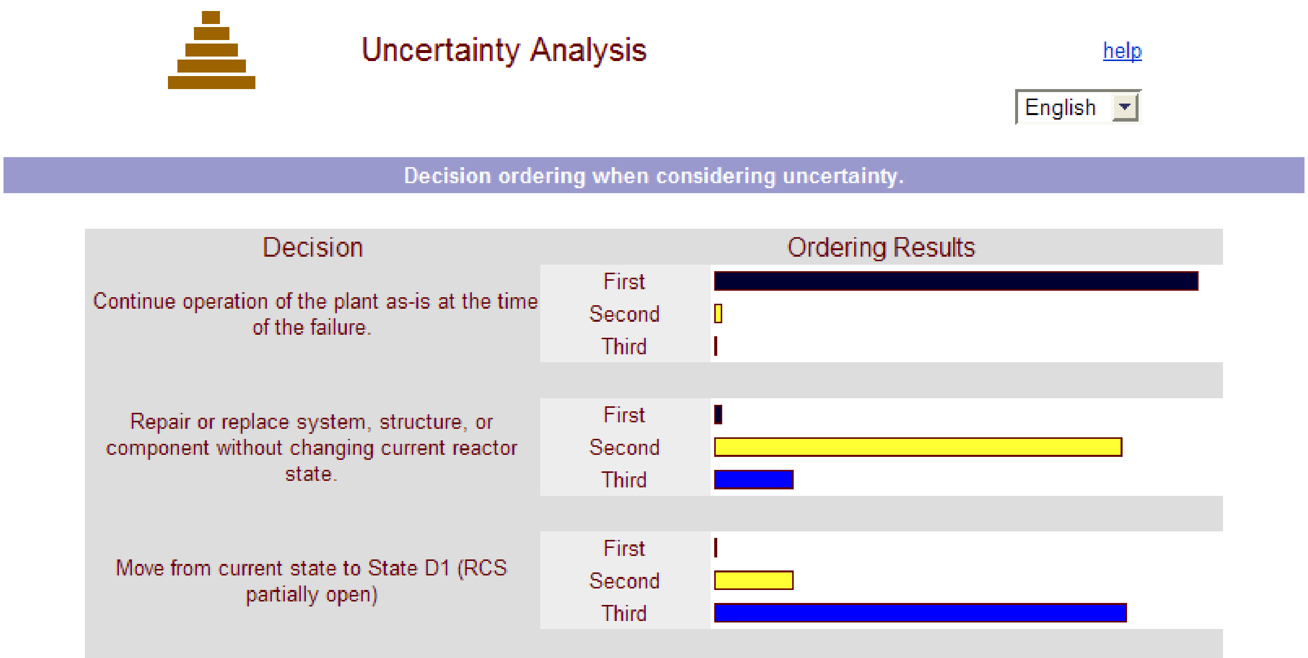

Figure 63. Uncertainty results for Case Study II.

"If all economists were laid end to end, they would not reach a conclusion."
— George Bernard Shaw

6 Conclusions

In this thesis, we have outlined the general framework behind formal decision making relevant to an incident management advisory system in nuclear power plants. Along the way, we faced several issues, including the determination of disutility functions, automation of the decision process, and the solution of time-dependent decision processes.

We were able to illustrate the complexities of our disutility functions by utilizing multi-dimensional consistency checks via fixed indifference levels. These "sanity" checks illustrated major shortcomings in the original framework that were ultimately rectified. But, a surprising result from our analysis was the difficulties encountered when using AHP for disutility determination. AHP is a technique that has seen increasing application in the field of decision analysis. But, we found several technical issues that limited its usefulness for our application, including:

- A tendency to produce excessive "risk prone" disutility functions for decision maker preferences.

- A flawed representation of extreme outcomes, specifically a poor fidelity in preserving the certainty equivalent at the upper end of our disutilities.

- An extreme sensitivity with regard to modifications of scales and matrix values at the lower end of the disutility curve.

- A limitation in transformations from AHP normalized weights to disutility values.

We initially used AHP to determine a disutility for our performance measures, measures that span several orders of magnitude. The developer of AHP, Saaty, has noted that AHP should *only* be used for "one order of magnitude" types of preference elicitation (Saaty, 1997). Consequently, to expand the use of AHP, Saaty has recommended a "pivot" approach whereby many smaller AHP regions are joined to construct one large AHP region, thereby avoiding the "order of magnitude" issue (Saaty and Vargas, 2000). But, from our analysis, we determined that this approach is flawed due to the nature of the AHP approach over-emphasizing the initial portion of the performance measure scale. We showed that by adjusting a single initial region, one could come up with dramatically different AHP-based disutility results. Since the number and magnitude of the scales for a disutility are arbitrary (for example, with a cardinal-based measure such as cost, one could have many intervals), the use of AHP is not consistent or robust for these types of applications.

For problems that involve (a) comparison outcomes that are within an order of magnitude and (b) have outcomes that are not arbitrary, then AHP may be useful. For example, in the evaluation of our value tree, the outcomes were measures like cost, worker safety, radiological dose, plant safety, and external attention. When we use AHP to determine the value tree weights, we are comparing one entity (the worth of cost) against another entity (the worth of worker safety). These value tree measures are within an order of magnitude from one another (one is not 100 or 1,000 times more important than the other in the context of incident management). Further, they do not have arbitrary outcomes (for example, on the cost measure, the focus is "cost" as an impact, not on any one particular value for cost). Consequently, we did use AHP to assist in the determination of value tree weights. But, AHP poised numerous challenges for the disutility determination.

We surmounted the problems of AHP-derived disutilities by developing a unique approach that we call "measurable equivalence." The measurable equivalence method uses facts, rather than judgement, to adjust disutilities. There are two guiding principles of this method. First, we ensure that the worst case for each performance measure has approximately the same level of "consequence." Second, where possible, the performance measure indifference points are constrained by actual measurable equalities, where a "measurable equality" is two performance measure outcomes that would be equivalent (e.g., a radiological dose of 7 Sv represents a fatality). This second principle is used to bring real data into the decision process while simultaneously reducing the subjectivity when utilizing preference information. Combined with the measurable equivalence approach, we utilized a cost disutility function based upon lottery equivalence input from our decision makers. The disutility functions embodied in the incident management advisor prototype reflect the measurable equivalence approach. Consistency checks that were preformed using these disutilities showed very good agreement with our expected results.

A second primary challenge we faced during this work was the integration and automation of the decision analysis theory and analysis heuristics, leading to a decision advisor prototype tool. In order to realize the successful implementation of this tool, we had to first decide on a programmatic framework for the software development – the result of this decision was to base the prototype on a web browser. Developing software for a browser causes some difficulties since a browser is stateless and HTML has a limited set of on-screen controls. But, the benefits of a familiar user interface, a robust open source application language (PHP version 4), and the potential for Internet collaboration amongst the research team outweighed the development complications. We were able to abstract the theory (Section 3) and methods (Section 4) in order to automate the incident decision process. Included in this process are:

- Decision maker preferences via the value tree (and corresponding weights) and the measurable equivalence disutilities.

- PRA models through the use of initiator upsets and safety system responses.

- Economic, worker safety, and radiological dose deterministic modeling.

- Decision alternatives related to incident management.

- Plant operational state determination.

- Sensitivity and uncertainty treatment.

Two major portions of the advisor prototype framework are the ideas of influenced-based information storage (via a knowledge base) and enveloping the entire decision process in a simulation representation.

In order to assist a user confronted with complex nuclear power plant decisions, we forced the prototype to rely on an underlying knowledge base. Rather than just having tables of information (e.g., lists of component names and descriptions) like one might find in a traditional database, we structured the database to allow us to capture decision-specific influences. For example, an auxiliary feedwater pump is a backup safety system that, if it fails, can generally be repaired while the plant stays at power. Thus, in our knowledge base, we defined a relationship between the decision node "repair at power" and the auxiliary feedwater pump. The positive outcome of this encoding of knowledge is that for any incident which involves a feedwater pump, the prototype will automatically allow the "repair at power" decision to enter into the final decision analysis model. While this example may seem obvious to individuals operating a nuclear power plant, this concept of influences between object nodes in the knowledge base allows us to define complex relationships that may otherwise be missed. The benefit of this relationship idea was revealed to us early into our research when we compared two decision models, one developed by us and one by our plant decision makers. The decision maker's model did *not* include the potential of causing an plant upset during repair (of a failed instrumentation module), even though the nuclear industry has seen these types of events. Consequently, by embedding operational information into the knowledge base, we make the resulting decision model robust and lessen the cognitive burden on the user. Further, as time passes and additional information is entered into the knowledge base, the decision advisor system becomes a mechanism of capturing operational history and institutional wisdom.

Since decision processes involve stochastic outcomes, we evaluated both traditional static and dynamic models. For treatment of the dynamic models, we utilized a version of the Metropolis simulation algorithm, but we extended the routine to encompass decision processes for nuclear power plants. We then ran this simulation approach and compared the results again analytic (exact) and approximate (static) model solutions. Within our decision framework, we are concerned with plant upsets – via initiating events – and the reliability of safety systems. We found that in some cases, static models will adequately represent the decision model. But, in other cases, static decision trees and fault trees provide only a rough approximation to the exact answer. Note though that static models may be solved very quickly, while simulation-based approaches take much longer. Consequently, we structured the decision advisor prototype such that a two-level analysis approach is possible. A static model can be used to provide decision advice immediately while the prototype continues to process the simulation model.

A variety of secondary accomplishments were realized during the research behind this report. Our treatment of epistemic uncertainty was comprehensive and included portions of the decision model such as preference functions and the value tree weights. The discussion of the potential modifications and limitations when using a PRA for decision making points out areas of concern. Since we do use a variety of PRA information as part of the decision advisor prototype, we developed a XML schema specific to this information in order to facilitate the transfer and manipulation of data structures found in a typical nuclear power plant PSA. We offer this format to the analysis community at large with the hope that it will encourage the exchange of reliability and safety information.

The decision advisor prototype that was constructed is considered to be Version 1 (the full source code is shown in Appendix D). The knowledge base that is behind the prototype was built such that, over time, plant-specific information that impacts decision making can be added as a permanent part of the advisor. Also, since the software uses a modularized framework, extensions to the prototype can be made as the need arises.

This thesis provides the script for a decision making regimen, a play that casts process models; applicable, informed decision makers; and formal decision-making technologies together with the goal of assisting, not replacing, human judgment. While decisions take place on the stage of uncertainty, it is important to remember that decision science provides a solid foundation for the framework described within this document. But, like any good story, our script has a beginning, middle, and (finally) an end. Alas, as the curtain falls, we have reached the end.

7 References

Accorsi, Roberto, George E. Apostolakis, and Enrico Zio, 1999, "Prioritizing Stakeholder Concerns in Environmental Risk Management," *Journal of Risk Research*, **2** (1), 11-29.

Ang, Alfredo H-S. and Wilson H. Tang, 1975, *Probability Concepts in Engineering Planning and Design, Vol. 1 Basic Principles*, John Wiley and Sons, pp. 11.

Aponte, Juan B. and Herbert S. Denenberg, 1968, "A New Concept of the Economics of Life Value and the Human Life Value: A Rationale for Term Insurance as the Cornerstone of Insurance Marketing," The Journal of Risk and Insurance, Vol. 35, Issue 3, September, 337-356.

Apostolakis, George E., and T. L. Chu, 1980, "The Unavailability of Systems under Period Test and Maintenance," *Nuclear Technology*, Vol. 50, Mid-August, pp. 5-15.

Apostolakis, George E., 1995, "A Commentary on Model Uncertainty," in: *Proceedings of Workshop on Model Uncertainty,"* A. Mosleh, N. Siu, C. Smidts, and C. Lui, Eds., Center for Reliability Engineering, University of Maryland, College Park, MD (also published as Report NUREG/CP-0138, US Nuclear Regulatory Commission, Washington, DC, 1994).

Apostolakis, George E., 2000, "Apostolakis: On PRA," *Nuclear News*, March, pp. 27-31

Atwood, Corwin L., 1996, "Constrained Noninformative Priors in Risk Assessment," *Reliability Engineering and System Safety*, **53**, 37-46.

Au, Siu-Kui and James L. Beck, 2001, "Estimation of Small Failure Probabilities in High Dimensions by Subset Simulation," *Probabilistic Engineering Mechanics*, **16**, pp. 263-277.

Bain, Lee J. and Max Engelhardt, 1992, *Introduction to Probability and Mathematical Statistics*, Duxbury Press, p. 201.

Bayes, Thomas, 1763, "An Essay Towards Solving a Problem in the Doctrine of Chances," *Philosophical Transactions of the Royal Society of London,* **53**, 370-418.

Bernstein, Peter L. 1996, *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons.

Blomquist, Glenn C., 2001, "Economics of the Value of Life," *International Encyclopedia of the Social and Behavioral Sciences*, Pergamon Press.

Brewer, Duncan and Ken S. Canady, 1999, "Probabilistic Safety Assessment Support for the Maintenance Rule at Duke Power Company," *Reliability Engineering and System Safety*, **63**, pp. 243-249.

Burke, Richard P., David C. Aldrich, Norman C. Rasmussen, 1984, *Economic Risks of Nuclear Power Reactor Accidents*, NUREG/CR-3673.

Call, Hollis J. and William A. Miller, 1990, "A Comparison of Approaches and Implementations for Automating Decision Analysis," *Reliability Engineering and System Safety*, **30**, pp. 115-162.

Cardano, Gerolamo, 1663, *Liber de ludo aleae (The Book on Games of Chance)*, translated by Sydney H. Gould, 1953, Princeton University Press, New Jersey.

Casti, John L., 1990, *Searching for Certainty*, William Morrow and Company, New York.

Claiborne, E. et al., 1989, *Generic Cost Estimates*, NUREG/CR-4627, Revision 1.

Catton, I., and W.E. Kastenberg, 1998, "Reactor Cavity Flooding as an Accident Management Strategy," *Reliability Engineering and System Safety*, **62**, pp. 59-70.

Cember, Herman, 1992, *Introduction to Health Physics*, McGraw-Hill.

Cheok, Michael C., Gareth W. Parry, and Richard R. Sherry, 1998, “Use of Importance Measures in Risk-Informed Regulatory Applications,” *Reliability Engineering and System Safety*, **60**, pp. 213-226.

Clemen, Robert T. 1996, *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press, 2nd edition.

Collier, Courtland A. and William B. Ledbetter, 1982, *Engineering Cost Analysis*, Harper & Row Publishers, New York.

Conlisk, John, 1989, “Three Variants on the Allais Example,” *The American Economic Review*, Vol. 79, No. 3., June, pp. 392-407.

Eerola, Mervi, 1994, *Probabilistic Causality in Longitudinal Studies*, Springer-Verlag, New York.

Edwards, W. 1954, “The Theory of Decision Making,” *Psychological Bulletin*, 51, 380-417.

Eide, Steve A. and Thomas E. Wierman, 2002, “Environment, Safety, And Health Risk Assessment Program (ESHRAP),” *PSAM 6, 6th International Conference on Probabilistic Safety Assessment and Management Conference Proceedings*, Elsevier Science.

Escobar, M. T. and J. M. Moreno-Jiménez, 2000, "Reciprocal Distributions in the Analytic Hierarchy Process," *European Journal of Operational Research*, **123**, pp. 154-174.

Forster, Malcolm R., 2001, "The New Science of Simplicity," *Simplicity, Inference and Modeling*, A. Zellner, H. A. Keuzenkamp, and M. McAleer (eds.), Cambridge University Press, pp. 83-119.

François, Philippe, 1998, "Probabilistic Analysis of the Potential Consequences of Incidents in Nuclear Power Plants: Methods, Lessons and Future Prospects," *Proceedings of the European Conference on Safety and Reliability – ESREL '98*, Edited by Stian Lydersen, Geir K. Hansen, and Helge A. Sandtorv. Published by A. A. Balkema, Netherlands, June.

Gertman, David I. et al., 1989, *Nuclear Computerized Library for Assessing Reactor Reliability (NUCLARR) Volume 1: Summary Description*, NUREG/CR-4639, February.

Goel, L. and Shuhua Ren, 1999, "Comprehensive Reliability Assessment of Substation Originated Outages Using the Monte Carlo Simulation Technique," *Electric Power Systems Research*, **52**, pp. 51-159.

Halmos, P. R., 1944, "The Foundations of Probability," *American Mathematical Monthly*, Vol. 51, Issue 9, 493-510.

Heiner, Ronald A., 1983, "The Origin of Predictable Behavior," *The American Economic Review*, Vol. 73, No. 4, September, pp. 560-595.

Henley, Ernest J. and Hiromitsu Kumamoto, 1981, *Reliability Engineering and Risk Assessment*, Prentice-Hall.

Hofflander, Alfred E., 1966, “The Human Life Value: An Historical Perspective,” *The Journal of Risk and Insurance*, Vol. 33, Issue 3 (Sept. 1966), 381-391.

Howard, Ronald A., 1971, *Dynamic Probabilistic Systems, Vol. I: Markov Models and Vol. II: Semimarkov and Decision Processes,* John Wiley & Sons, New York.

Hughes, Warren R., 1986, “Deriving Utilities Using the Analytic Hierachy Process,” *Socio-Economic Planning Science*, Vol. 20, No. 6, p. 394.

Jae, Moosung and George E. Apostolakis, 1992, “The Use of Influence Diagrams for Evaluating Severe Accident Management Strategies,” *Nuclear Technologies*, Vol. 99, August, pp. 142-157.

Kahneman, Daniel and Amos Tversky, 1979, “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica*, Vol. 47, No. 2, March, pp.279.

Keeney, Ralph L. and Howard Raiffa, 1993, *Decisions with Multiple Objectives*, Cambridge University Press.

Keeney, Ralph L., 1997, “Estimating Fatalities Induced by the Economic Costs of Regulations,” *The Journal of Risk and Uncertainty*, Vol. 14, No. 5, pp. 23.

Knight, Frank H., 1921, *Risk, Uncertainty, and Profit*, Houghton Mifflin Company, Library of Economics and Liberty, http://www.econlib.org/library/Knight/knRUP1.html.

Kovan, Dick, 2002, “Outage strategy well covered at Forsmark,” *Nuclear News*, pp. 32-39, April.

Kim, Chul and Hoe Kyung Lee, 1992, “A Monte Carlo Simulation Algorithm for Finding MTBF,” *IEEE Transactions on Reliability*, Vol. 41, No. 2, pp. 193-195, June.

Lambert, Howard E., 1975, *Fault Trees for Decision Making in Systems Analysis*, UCRL-51829, October.

Lambert, Howard E., 1975, “Measure of Importance of Events and Cut Sets in Fault Trees,” *Reliability and Fault Tree Analysis*, ed. R. E. Barlow, J. B. Fussell, and N. D. Singpurwalla, SIAM Press, Philadelphia, pp. 77-100.

Laplace, Pierre-Simon, 1995, *Essai Philosophique sur les Probabilités (Philosophical Essay on Probabilities)*, translated by Andrew I. Dale, Springer-Verlag, New York.

Lindstedt, Mats, Raimo Hämäläinen, Jyri Mustajoki, 2001, “Using Intervals for Global Sensitivity Analyses in Multiattribute Value Trees,” *Proceedings of the Fifteenth International Conference on Multiple Criteria Decision Making*, Murat Köksalan and Stanley Zionts (eds.), pp. 177-186.

Lopez, B. and F. W. Sciacca, 1990, *FORECAST, 2.1 User Manual*, Science and Engineering Associates, Inc., SEA Report No. 89-461-04-A:1.

Machina, Mark J., 1982, “’Expected Utility’ Analysis without the Independence Axiom,” *Econometrica*, Vol. 50, No. 2, pp. 277-324.

Machina, Mark J., 1987, “Decision-Making in the Presence of Risk,” *Science*, New Series, Vol. 236, No. 4801, pp. 537-543.

Magee, John, 1964, “Decision Trees for Decision Making,” *Harvard Business Review*, July/August.

Martz, Harry F. and Ray. A. Waller, 1991, *Bayesian Reliability Analysis*, Krieger Publishing Company, 2nd edition.

McCormick, Norman J., 1981, *Reliability and Risk Analysis*, Academic Press.

Merkhofer, Miley W. and Ralph L. Keeney, 1987, A Multiattribute Utility Analysis of Alternative Sites for the Disposal of Nuclear Waste, *Risk Analysis*, Vol. 7, No. 2, pp. 173-194.

Metropolis, Nicholas, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller, 1953, "Equations of State Calculations by Fast Computing Machines," *Journal of Chemical Physics*, Vol. 21, No. 6, June, pp. 1087-92.

Milici, A., J.-S. Wu, G. Apostolakis, K-F. Yau, 1995, *Accident Management Advisor System Development to handle Uncertain Sensor Inputs under Nuclear Plant Accident Conditions*, prepared for U.S. Department of Energy, Phase II Report.

Modarres, Mohammad, 1993, *What Every Engineer Should Know About Reliability and Risk Analysis*, Marcel Dekker.

Neumann, John von, and Oskar Morgenstern, 1944, *Theory of Games and Economic Behavior*, Princeton University Press, New Jersey.

Nuclear News, 2001, "Analysts hear overview of nuclear power's value," p. 15, March.

Papazoglou, Ioannis A. and Elias P. Gyftopoulos, 1977, "Markov Processes for Reliability Analyses of Large Systems," *IEEE Transactions on Reliability*, Vol. R-26, No. 3, August, pp. 232-237.

Persson, Ulf, Anna Norinder, Krister Hjalte, and Katarina Gralén, 2001, "The Value of a Statistical Life in Transport: Findings from a New Contingent Valuation Study in Sweden," *The Journal of Risk and Uncertainty,* Vol. 23, No. 2, 121-134.

Phelps, E. S. and R. A. Pollak, 1968, "On Second-best National Saving and Game-equilibrium Growth," *Review of Economic Studies*, Vol. 35, 189-199.

Pöyhönen, Mari and Raimo P. Hämäläinen, 2001, "On the Convergence of Multiattribute Weighting Methods," *European Journal of Operational Research*, Vol. 129, 569-585.

Psillakis, Zaharias, 1995, "A Simulation Algorithm for Computing Failure Probability of a Consequtive-k-out-of-r-from-n:F System," *IEEE Transactions on Reliability*, Vol. 44, No. 3, pp. 523-531, September.

Raiffa, Howard and Robert Schlaifer, 1961, *Applied Statistical Decision Theory,* Division of Research, Graduate School of Business Administration, Harvard University.

Raiffa, Howard, 1968, *Decision Analysis, Introductory Lectures on Choice Under Uncertainty*, Addison-Wesley, Reading, Massachusetts.

Ripley, Brian D., 1987, *Stochastic Simulation*, John Wiley & Sons, New York.

Rodgers, Shawn S. and Karl Fleming, 1998, "Property Damage Risk Assessment for Nuclear Power Plants," *Probabilistic Safety Assessment and Management, PSAM 4*, A. Mosleh and R.A. Bari (eds.), Springer-Verlag.

Saaty, Thomas, 1980, *The Analytic Hierarchy Process,* McGraw Hill.

Saaty, Thomas, 1997, "That is Not the Analytic Hierachy Process: What the AHP is and What It Is Not," *The Journal of Multi-Criteria Decision Analysis*, Vol. 6, pp. 324-335.

Saaty, Thomas L. and Luis G. Vargas, 2000, *Models, Methods, Concepts & Applications of the Analytic Hierarchy Process (International Series in Operations Research and Management Science, Volume 34)*, Kluwer Academic Publishers.

Salo, Ahti A. and Raimo P. Hämäläinen, 1997, "On the Measurement of Preferences in the Analytic Hierarchy Process," *The Journal of Multi-Criteria Decision Analysis*, Vol. 6, pp. 309-319.

Sciacca, Frank. W., 1989, *Generic Cost Estimates*, NUREG/CR-4627, Revision 2, U.S. Nuclear Regulatory Commission, Washington, D.C.

Siu, Nathan, 1994, "Risk Assessment for Dynamic Systems: An Overview," *Reliability Engineering and System Safety*, Vol. 43, pp. 43-73.

Smith, Curtis L., 1998, "Calculating Conditional Core Damage Probabilities for Nuclear Power Plant Operations," *Reliability Engineering and System Safety*, vol. 59, pp. 299-307.

Smith, Curtis L., James K. Knudsen, and Michael. B. Calley, 1999, "Calculating and addressing uncertainty for risk-based allowable outage times," *Reliability Engineering and System Safety*, Vol. 66, pp. 41-47.

Smith, Curtis L., John A. Schroeder, Scott T. Beck, and James K. Knudsen, 2002, "Modeling Power Non-Recovery using the SAPHIRE Risk Assessment Software," *PSAM 6, 6th International Conference on Probabilistic Safety Assessment and Management Conference Proceedings*, Elsevier Science.

Swain, A. D. and H. E. Guttmann, 1983, *Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications - Final Report*, NUREG/CR-1278-F, August.

Swain, A. D., 1987, *Accident Sequence Evaluation Program Human Reliability Analysis Procedure*, NUREG/CR-4772, February.

Todhunter, Isaac, 1865, *A History of the Mathematical Theory of Probability from the Time of Pascal to that of Laplace*, Thoemmes Press, Bristol, England (reprint 2001).

Weil, Rick. and George E. Apostolakis, 2001, "A Methodology for the Prioritization of Operating Experience in Nuclear Power Plants," *Reliability Engineering and System Safety*, Vol. 74.

Winkler, Robert L., 1972, *An Introduction to Bayesian Inference and Decision*, Holt, Rinehart, and Winston, New York, 1972.

Wright, George and Paul Goodwin, 1999, "Rethinking Value Elicitation for Personal Consequential Decisions," *Journal of Multi-Criteria Decision Analysis*, Vol. 8, pp. 3-10.

Wu-Chien, Jya Syin and George E. Apostolakis, 1983, "Prospect Theory and Limit Lines," *Risk Analysis*, Vol. 3, No. 3, pp. 181-188.

U.S. Government, 1990, *Report to the Congress from the Presidential Commission on Catastrophic Nuclear Accidents*, , U.S. Government Printing Office, Washington, D.C.

U.S. Government, 2002, *Economic Report of the President*, U.S. Government Printing Office, Washington, D.C.

U.S. NRC, 1997, *Regulatory Analysis Technical Evaluation Handbook*, NUREG/BR-0184, Office of Nuclear Regulatory Research, January.

U.S. NRC, 2000, *Risk-Informing Special Treatment Requirements*, 10 CFR Parts 21, 50, 52, 54 and 100, Rulemaking Information Notice (RIN) 3150-AG42.